Algorithms for programmers
ideas and source code

This document is work in progress: read the "important remarks" near the beginning

Jörg Arndt
arndt@jjj.de

Draft version[1] of 2008-April-17

[1]The latest version and the accompanying software is online at http://www.jjj.de/fxt/.

Contents

Important remarks about this document **xi**

I Low level algorithms 1

1 Bit wizardry **3**
1.1 Trivia . . . 3
1.2 Operations on individual bits . . . 8
1.3 Operations on low bits or blocks of a word . . . 9
1.4 Extraction of ones, zeros, or blocks near transitions . . . 12
1.5 Computing the index of a single set bit . . . 14
1.6 Operations on high bits or blocks of a word . . . 16
1.7 Functions related to the base-2 logarithm . . . 18
1.8 Counting bits and blocks of a word . . . 19
1.9 Counting bits of many words . . . 22
1.10 Bit set lookup . . . 24
1.11 Avoiding branches . . . 24
1.12 Bit-wise rotation of a word . . . 28
1.13 Functions related to bit-wise rotation and binary necklaces . . . 29
1.14 Reversing the bits of a word . . . 34
1.15 Bit-wise zip . . . 39
1.16 Gray code and parity . . . 40
1.17 Bit sequency . . . 46
1.18 Powers of the Gray code . . . 48
1.19 Invertible transforms on words . . . 49
1.20 Moves of the Hilbert curve . . . 56
1.21 The Z-order . . . 57
1.22 Scanning for zero bytes . . . 58
1.23 2-adic inverse and square root . . . 59
1.24 Radix −2 representation . . . 61
1.25 A sparse signed binary representation . . . 64
1.26 Generating bit combinations . . . 66
1.27 Generating bit subsets of a given word . . . 70
1.28 Binary words as subsets in lexicographic order . . . 71
1.29 Minimal-change bit combinations . . . 76
1.30 Fibonacci words . . . 78
1.31 Binary words and parentheses strings * . . . 81
1.32 Error detection by hashing: the CRC . . . 84
1.33 Permutations via primitives . . . 88
1.34 CPU instructions often missed . . . 91

2 Permutations **93**

2.1 The revbin permutation . . . 93
2.2 The radix permutation . . . 97
2.3 In-place matrix transposition . . . 98
2.4 Revbin permutation and matrix transposition * . . . 99
2.5 The zip permutation . . . 101
2.6 The reversed zip permutation . . . 103
2.7 The XOR permutation . . . 104
2.8 The Gray code permutation . . . 105
2.9 The reversed Gray code permutation . . . 109
2.10 Decomposing permutations * . . . 110
2.11 General permutations and their operations . . . 112

3 Sorting and searching 123
3.1 Sorting . . . 123
3.2 Binary search . . . 125
3.3 Index sorting . . . 126
3.4 Pointer sorting . . . 128
3.5 Sorting by a supplied comparison function . . . 129
3.6 Determination of unique elements . . . 132
3.7 Unique elements with inexact types . . . 133
3.8 Determination of equivalence classes . . . 135
3.9 Determination of monotonicity and convexity * . . . 139
3.10 Heapsort . . . 142
3.11 Counting sort and radix sort . . . 142
3.12 Searching in unsorted arrays . . . 145

4 Data structures 149
4.1 Stack (LIFO) . . . 149
4.2 Ring buffer . . . 151
4.3 Queue (FIFO) . . . 152
4.4 Deque (double-ended queue) . . . 154
4.5 Heap and priority queue . . . 157
4.6 Bit-array . . . 160
4.7 Finite-state machines . . . 162
4.8 Emulation of coroutines . . . 164

II Combinatorial generation 169

5 Conventions and considerations 171
5.1 About representations and orders . . . 171
5.2 Ranking, unranking, and counting . . . 172
5.3 Characteristics of the algorithms . . . 172
5.4 Optimization techniques . . . 172
5.5 Remarks about the C++ implementations . . . 174

6 Combinations 175
6.1 Lexicographic and co-lexicographic order . . . 176
6.2 Order by prefix shifts (cool-lex) . . . 179
6.3 Minimal-change order . . . 181
6.4 The Eades-McKay strong minimal-change order . . . 182
6.5 Two-close orderings via endo/enup moves . . . 185
6.6 Recursive generation of certain orderings . . . 190

7 Compositions **193**
7.1 Co-lexicographic order . . . 193
7.2 Co-lexicographic order for compositions into exactly k parts . . . 195
7.3 Compositions and combinations . . . 197
7.4 Minimal-change orders . . . 198

8 Subsets **201**
8.1 Lexicographic order . . . 201
8.2 Minimal-change order . . . 203
8.3 Ordering with De Bruijn sequences . . . 207
8.4 Shifts-order for subsets . . . 208
8.5 k-subsets where k lies in a given range . . . 210

9 Mixed radix numbers **219**
9.1 Counting order . . . 219
9.2 Gray code order . . . 222
9.3 gslex order . . . 225
9.4 endo order . . . 227
9.5 Gray code for endo order . . . 230

10 Permutations **233**
10.1 Lexicographic order . . . 233
10.2 Co-lexicographic order . . . 235
10.3 Factorial representations of permutations . . . 236
10.4 An order from reversing prefixes . . . 245
10.5 Minimal-change order (Heap's algorithm) . . . 248
10.6 Lipski's Minimal-change orders . . . 250
10.7 Strong minimal-change order (Trotter's algorithm) . . . 254
10.8 Minimal-change orders from factorial numbers . . . 258
10.9 Orders where the smallest element always moves right . . . 265
10.10 Single track orders . . . 269
10.11 Star-transposition order . . . 274
10.12 Derangement order . . . 275
10.13 Recursive algorithm for cyclic permutations . . . 278
10.14 Minimal-change order for cyclic permutations . . . 280
10.15 Permutations with special properties . . . 282

11 Subsets and permutations of a multiset **289**
11.1 Subsets of a multiset . . . 289
11.2 Permutations of a multiset . . . 290

12 Gray codes for strings with restrictions **295**
12.1 Fibonacci words . . . 296
12.2 Generalized Fibonacci words . . . 298
12.3 Digit x followed by at least x zeros . . . 301
12.4 Generalized Pell words . . . 302
12.5 Sparse signed binary words . . . 304
12.6 Strings with no two successive nonzero digits . . . 306
12.7 Strings with no two successive zeros . . . 308
12.8 Binary strings without substrings 1x1 or 1xy1 . . . 309

13 Parenthesis strings **313**
13.1 Co-lexicographic order . . . 313
13.2 Gray code via restricted growth strings . . . 315

13.3 The number of parenthesis strings: Catalan numbers . . . 320
13.4 Increment-i RGS and k-ary trees . . . 321

14 Integer partitions **327**
14.1 Recursive solution of a generalized problem . . . 327
14.2 Iterative algorithm . . . 329
14.3 Partitions into m parts . . . 331
14.4 The number of integer partitions . . . 332

15 Set partitions **335**
15.1 The number of set partitions: Stirling set numbers and Bell numbers . . . 336
15.2 Generation in minimal-change order . . . 338
15.3 Restricted growth strings . . . 341

16 A string substitution engine **349**

17 Necklaces and Lyndon words **353**
17.1 Generating all necklaces . . . 354
17.2 The number of binary necklaces . . . 361
17.3 The number of binary necklaces with fixed content . . . 362

18 Hadamard and conference matrices **365**
18.1 Hadamard matrices via LFSR . . . 365
18.2 Hadamard matrices via conference matrices . . . 367
18.3 Conference matrices via finite fields . . . 369

19 Searching paths in directed graphs **373**
19.1 Representation of digraphs . . . 374
19.2 Searching full paths . . . 375
19.3 Conditional search . . . 380
19.4 Edge sorting and lucky paths . . . 384
19.5 Gray codes for Lyndon words . . . 385

III Fast orthogonal transforms 391

20 The Fourier transform **393**
20.1 The discrete Fourier transform . . . 393
20.2 Summary of definitions of Fourier transforms * . . . 394
20.3 Radix-2 FFT algorithms . . . 396
20.4 Saving trigonometric computations . . . 401
20.5 Higher radix FFT algorithms . . . 403
20.6 Split-radix Fourier transforms . . . 411
20.7 Symmetries of the Fourier transform . . . 414
20.8 Inverse FFT for free . . . 415
20.9 Real valued Fourier transforms . . . 416
20.10 Multidimensional Fourier transforms . . . 423
20.11 The matrix Fourier algorithm (MFA) . . . 424

21 Algorithms for fast convolution **427**
21.1 Convolution . . . 427
21.2 Correlation . . . 432
21.3 Weighted Fourier transforms and convolutions . . . 435
21.4 Convolution using the MFA . . . 437
21.5 The z-transform (ZT) . . . 440

21.6 Prime length FFTs . . . 443

22 The Walsh transform and its relatives **447**
22.1 The Walsh transform: Walsh-Kronecker basis . . . 447
22.2 Eigenvectors of the Walsh transform * . . . 450
22.3 The Kronecker product . . . 451
22.4 A variant of the Walsh transform * . . . 453
22.5 Higher radix Walsh transforms . . . 455
22.6 Localized Walsh transforms . . . 459
22.7 Dyadic (XOR) convolution . . . 463
22.8 The Walsh transform: Walsh-Paley basis . . . 465
22.9 Sequency ordered Walsh transforms . . . 466
22.10 Slant transform . . . 472
22.11 Arithmetic transform . . . 473
22.12 Reed-Muller transform . . . 476
22.13 The OR-convolution, and the AND-convolution . . . 479

23 The Haar transform **483**
23.1 The 'standard' Haar transform . . . 483
23.2 In-place Haar transform . . . 485
23.3 Non-normalized Haar transforms . . . 487
23.4 Transposed Haar transforms . . . 489
23.5 The reversed Haar transform . . . 491
23.6 Relations between Walsh and Haar transforms . . . 493
23.7 Nonstandard splitting schemes * . . . 496

24 The Hartley transform **501**
24.1 Definition and symmetries . . . 501
24.2 Radix-2 FHT algorithms . . . 502
24.3 Complex FT by HT . . . 507
24.4 Complex FT by complex HT and vice versa . . . 508
24.5 Real FT by HT and vice versa . . . 509
24.6 Higher radix FHT algorithms . . . 510
24.7 Convolution via FHT . . . 511
24.8 Negacyclic convolution via FHT . . . 515
24.9 Localized FHT algorithms . . . 515
24.10 Two-dimensional FHTs . . . 517
24.11 Discrete cosine transform (DCT) by HT . . . 518
24.12 Discrete sine transform (DST) by DCT . . . 519
24.13 Automatic generation of transform code . . . 520
24.14 Eigenvectors of the Fourier and Hartley transform * . . . 522

25 Number theoretic transforms (NTTs) **525**
25.1 Prime moduli for NTTs . . . 525
25.2 Implementation of NTTs . . . 527
25.3 Convolution with NTTs . . . 532

26 Fast wavelet transforms **533**
26.1 Wavelet filters . . . 533
26.2 Implementation . . . 535
26.3 Moment conditions . . . 536

IV Fast arithmetic 539

27 Fast multiplication and exponentiation 541
27.1 Asymptotics of algorithms . . . 541
27.2 Splitting schemes for multiplication . . . 542
27.3 Fast multiplication via FFT . . . 550
27.4 Radix/precision considerations with FFT multiplication . . . 552
27.5 The sum-of-digits test . . . 554
27.6 Binary exponentiation . . . 555

28 Root extraction 559
28.1 Division, square root and cube root . . . 559
28.2 Root extraction for rationals . . . 562
28.3 Divisionless iterations for the inverse a-th root . . . 565
28.4 Initial approximations for iterations . . . 568
28.5 Some applications of the matrix square root . . . 569
28.6 Goldschmidt's algorithm . . . 574
28.7 Products for the a-th root . . . 577
28.8 Divisionless iterations for polynomial roots . . . 579

29 Iterations for the inversion of a function 581
29.1 Iterations and their rate of convergence . . . 581
29.2 Schröder's formula . . . 582
29.3 Schröder's formula and series reversion . . . 583
29.4 Householder's formula . . . 586
29.5 Dealing with multiple roots . . . 587
29.6 More iterations . . . 589
29.7 Convergence improvement by the delta squared process . . . 591

30 The arithmetic-geometric mean (AGM) 593
30.1 The AGM . . . 593
30.2 The elliptic functions K and E . . . 595
30.3 AGM-type algorithms for hypergeometric functions . . . 598
30.4 Computation of π . . . 602
30.5 Arctangent relations for π * . . . 610

31 Logarithm and exponential function 617
31.1 Logarithm . . . 617
31.2 Exponential function . . . 623
31.3 Logarithm and exponential function of power series . . . 626
31.4 Simultaneous computation of logarithms of small primes . . . 628

32 Numerical evaluation of power series 631
32.1 The binary splitting algorithm for rational series . . . 631
32.2 Rectangular schemes for evaluation of power series . . . 637
32.3 The magic sumalt algorithm for alternating series . . . 641

33 Computing the elementary functions with limited resources 645
33.1 Shift-and-add algorithms for $\log_b(x)$ and b^x . . . 645
33.2 CORDIC algorithms . . . 650

34 Recurrences and Chebyshev polynomials 657
34.1 Recurrences . . . 657
34.2 Chebyshev polynomials . . . 667

35 Cyclotomic polynomials, Hypergeometric functions, and continued fractions 677
35.1 Cylotomic polynomials, Möbius inversion, Lambert series . . . 677
35.2 Hypergeometric functions . . . 685
35.3 Continued fractions . . . 702

36 Synthetic Iterations * 713
36.1 A variation of the iteration for the inverse . . . 713
36.2 An iteration related to the Thue constant . . . 717
36.3 An iteration related to the Golay-Rudin-Shapiro sequence . . . 718
36.4 Iterations related to the ruler function . . . 720
36.5 An iteration related to the period-doubling sequence . . . 722
36.6 An iteration from substitution rules with sign . . . 726
36.7 Iterations related to the sum of digits . . . 726
36.8 Iterations related to the binary Gray code . . . 728
36.9 A function that encodes the Hilbert curve . . . 734
36.10 Sparse variants of the inverse . . . 737
36.11 An iteration related to the Fibonacci numbers . . . 740
36.12 Iterations related to the Pell numbers . . . 744

V Algorithms for finite fields 751

37 Modular arithmetic and some number theory 753
37.1 Implementation of the arithmetic operations . . . 753
37.2 Modular reduction with structured primes . . . 759
37.3 The sieve of Eratosthenes . . . 759
37.4 The order of an element . . . 761
37.5 Prime modulus: the field $\mathbb{Z}/p\mathbb{Z} = \mathbb{F}_p = \mathrm{GF}(p)$ . . . 762
37.6 Composite modulus: the ring $\mathbb{Z}/m\mathbb{Z}$ . . . 763
37.7 The Chinese Remainder Theorem (CRT) . . . 769
37.8 Quadratic residues . . . 771
37.9 Computation of a square root modulo m . . . 773
37.10 The Rabin-Miller test for compositeness . . . 775
37.11 Proving primality . . . 781
37.12 Complex moduli: $\mathrm{GF}(p^2)$ . . . 793
37.13 Solving the Pell equation . . . 801
37.14 Multigrades * . . . 804
37.15 Properties of the convergents of $\sqrt{2}$ * . . . 805
37.16 Multiplication of hypercomplex numbers * . . . 810

38 Binary polynomials 815
38.1 The basic arithmetical operations . . . 815
38.2 Multiplication for polynomials of high degree . . . 821
38.3 Modular arithmetic with binary polynomials . . . 827
38.4 Irreducible polynomials . . . 831
38.5 Primitive polynomials . . . 834
38.6 The number of irreducible and primitive polynomials . . . 837
38.7 Transformations that preserve irreducibility . . . 838
38.8 Self-reciprocal polynomials . . . 840
38.9 Irreducible and primitive polynomials of special forms * . . . 842
38.10 Generating irreducible polynomials from Lyndon words . . . 849
38.11 Irreducible and cyclotomic polynomials * . . . 850
38.12 Factorization of binary polynomials . . . 851

39 Shift registers **857**
39.1 Linear feedback shift registers (LFSR) . . . 857
39.2 Galois and Fibonacci setup . . . 860
39.3 Generating all revbin pairs . . . 861
39.4 The number of m-sequences and De Bruijn sequences . . . 862
39.5 Auto correlation of m-sequences . . . 863
39.6 Feedback carry shift register (FCSR) . . . 864
39.7 Linear hybrid cellular automata (LHCA) . . . 866
39.8 Additive linear hybrid cellular automata . . . 871

40 Binary finite fields: GF(2^n) **875**
40.1 Arithmetic and basic properties . . . 875
40.2 Minimal polynomials . . . 882
40.3 Fast computation of the trace vector . . . 884
40.4 Solving quadratic equations . . . 886
40.5 Representation by matrices * . . . 888
40.6 Representation by normal bases . . . 890
40.7 Conversion between normal and polynomial representation . . . 899
40.8 Optimal normal bases (ONB) . . . 900
40.9 Gaussian normal bases . . . 903

A The electronic version of the book **911**

B Machine used for benchmarking **913**

C The pseudo language Sprache **915**

D The pari/gp language **917**

Bibliography **925**

Index **943**

Important remarks about this document

This is a draft of a book about selected algorithms. The audience in mind are programmers who are interested in the treated algorithms and actually want to create and understand working and reasonably optimized code.

The style varies somewhat which I do not consider bad per se: While some topics (as fast Fourier transforms) need a clear and explicit introduction others (like the bit wizardry chapter) seem to be best presented by basically showing the code with just a few comments.

The pseudo language Sprache is used when I see a clear advantage to do so, mainly when the corresponding C++ does not appear to be self explanatory. Larger pieces of code are presented in C++. C programmers do not need to be shocked by the '++' as only a rather minimal set of the C++ features is used. Some of the code, especially in part 3 (Arithmetical algorithms), is given in the pari/gp language as the use of other languages would likely bury the idea in technicalities.

A printable version of this book will always stay online for free download. The referenced sources are online as part of FXT (fast transforms and low level routines [19]) and hfloat (high precision floating point algorithms [20]).

The reader is welcome to criticize and suggest improvements. Please name the draft version (date) with your feedback! **This version is of 2008-April-17.** Note that you can copy and paste from the PDF and DVI versions. Thanks go to those[1] who helped to improve this document so far!

In case you want to cite this document, please avoid referencing individual chapters or sections as their numbers (and titles) may change.

Enjoy reading!

Legal matters

This book is copyright © by its author, Jörg Arndt.
Redistributing or selling this book in printed or in electronic form is prohibited.
This book must not be mirrored on the Internet.
Using this book as promotional material is prohibited.
CiteSeer (`http://citeseer.ist.psu.edu/cs/`, and its mirrors) is allowed to keep a copy of this book in its database.

[1] in particular Igal Aharonovich, Max Alekseyev, Nathan Bullock, Dominique Delande, Torsten Finke, Sean Furlong, Almaz Gaifullin, Alexander Glyzov, Andreas Grünbacher, Christoph Haenel, Tony Hardie-Bick, Laszlo Hars, Jeff Hurchalla, Gideon Klimer, Richard B. Kreckel, Gál László, Dirk Lattermann, Avery Lee, Brent Lehman, Marc Lehmann, Paul C. Leopardi, John Lien, Mirko Liss, Johannes Middeke, Doug Moore, Andrew Morris, David Nalepa, Matthew Oliver, Mirosław Osys, Christoph Pacher, Scott Paine, Yves Paradis, Edith Parzefall, André Piotrowski, David García Quintas, Tony Reix, Johan Rönnblom, Thomas Schraitle, Clive Scott, Michael Somos, Michal Staruch, Ralf Stephan, Mikko Tommila, Michael Roby Wetherfield, Jim White, Vinnie Winkler, John Youngquist, Rui Zhang, and Paul Zimmermann.

"Why make things difficult, when it is possible to make them cryptic and totally illogical, with just a little bit more effort?"

– Aksel Peter Jørgensen

Part I

Low level algorithms

Chapter 1

Bit wizardry

We present low-level functions that operate on the bits of a binary word. It is often not obvious what these are good for and I do not attempt much to motivate why particular functions are presented. However, *if* you happen to have an application for a given routine you will love that it is there: the program using it may run significantly faster.

The C-type `unsigned long` is abbreviated as `ulong` as defined in [FXT: fxttypes.h]. It is assumed that `BITS_PER_LONG` reflects the size of an `unsigned long`. It is defined in [FXT: bits/bitsperlong.h] and (on sane architectures) equals the machine word size. That is, it equals 32 on 32-bit architectures and 64 on 64-bit machines. Further, the quantity `BYTES_PER_LONG` shall reflect the number of bytes in a machine word, that is, it equals `BITS_PER_LONG` divided by eight. For some functions it is assumed that `long` and `ulong` have the same number of bits.

The examples of assembler code are for the x86 and the AMD64 architecture. They should be simple enough to be understandable for readers who know assembler for any CPU.

1.1 Trivia

1.1.1 Little endian versus big endian

The order in which the bytes of an integer are stored in memory can start with the least significant byte (*little endian* machine) or with the most significant byte (*big endian* machine). The hexadecimal number `0x0D0C0B0A` will be stored in the following manner when memory addresses grow from left to right:

```
adr:   z   z+1   z+2  z+3
mem:   0D   0C   0B   0A   // big endian
mem:   0A   0B   0C   0D   // little endian
```

The difference is only visible when you cast pointers. Let `V` be the 32-bit integer with the value above. Then the result of `char c = *(char *)(&V);` will be `0x0A` (value modulo 256) on a little endian machine but `0x0D` (value divided by 2^{24}) on a big endian machine. Portable code that uses casts may need two versions, one for big endian, and one for small endian machines. Though friends of the big endian way sometimes refer to little endian as 'wrong endian', the desired result of the shown pointer cast is much more often the modulo operation.

1.1.2 Size of pointer is size of long

On sane architectures a pointer fits into a type `long` integer. When programming for a 32-bit architecture (where the size of `int` and `long` coincide) casting pointers to integers (and back) will work. The same

code *will* fail on 64-bit machines. If you have to cast pointers to an integer type, cast them to `long`.

1.1.3 Shifts and division

With two's complement arithmetic (that is: on likely every computer you'll ever touch) division and multiplication by powers of two is right and left shift, respectively. This is true for unsigned types and for multiplication (left shift) with signed types. Division with signed types rounds toward zero, as one would expect, but right shift is a division (by a power of two) that rounds to minus infinity:

```
int a = -1;
int c = a >> 1;   // c == -1
int d = a / 2;    // d ==  0
```

The compiler still uses a shift instruction for the division, but with a 'fix' for negative values:

```
  9:test.cc   @ int foo(int a)
 10:test.cc   @ {
285 0003 8B442410              movl 16(%esp),%eax  // move argument to %eax
 11:test.cc   @     int s = a >> 1;
289 0007 89C1                  movl %eax,%ecx
290 0009 D1F9                  sarl $1,%ecx
 12:test.cc   @     int d = a / 2;
293 000b 89C2                  movl %eax,%edx
294 000d C1EA1F                shrl $31,%edx  // fix: %edx=(%edx<0?1:0)
295 0010 01D0                  addl %edx,%eax // fix: add one if a<0
296 0012 D1F8                  sarl $1,%eax
```

For unsigned types the shift would suffice. One more reason to use unsigned types whenever possible.

The assembler listing was generated from C code via the following commands:

```
# create assembler code:
c++ -S -fverbose-asm -g -O2  test.cc -o test.s
# create asm interlaced with source lines:
as -alhnd test.s  > test.lst
```

There are two types of *right* shifts: a so-called logical and an arithmetical shift. The logical version (`shrl` in the above fragment) always fills the higher bits with zeros, corresponding to division[1] of unsigned types. The arithmetical shift (`sarl` in the above fragment) fills in ones or zeros, according to the most significant bit of the original word.

Computing remainders modulo a power of two with unsigned types is equivalent to a bit-and using a mask:

```
ulong a = b % 32;  // == b & (32-1)
```

All of the above is done by the compiler's optimization wherever possible.

Division by (compile time) constants can be replaced by multiplications and shift. The magic machinery inside the compiler does it for you. A division by the constant 10 is compiled to:

```
  5:test.cc   @ ulong foo(ulong a)
  6:test.cc   @ {
  7:test.cc   @     ulong b = a / 10;
290 0000 8B442404              movl 4(%esp),%eax
291 0004 F7250000              mull .LC33   // value == 0xcccccccd
292 000a 89D0                  movl %edx,%eax
293 000c C1E803                shrl $3,%eax
```

Thereby it is sometimes reasonable to have separate code branches with explicit special values. Similarly, for modulo computations with a constant modulus (using modulus 10,000):

```
  8:test.cc   @ ulong foo(ulong a)
  9:test.cc   @ {
 53 0000 8B4C2404              movl 4(%esp),%ecx
 10:test.cc   @     ulong b = a % 10000;
 57 0004 89C8                  movl %ecx,%eax
 58 0006 F7250000              mull .LC0   // value == 0xd1b71759
 59 000c 89D0                  movl %edx,%eax
 60 000e C1E80D                shrl $13,%eax
 61 0011 69C01027              imull $10000,%eax,%eax
```

[1] So you can think of it as 'unsigned arithmetical' shift.

```
62 0017 29C1                subl %eax,%ecx
63 0019 89C8                movl %ecx,%eax
```

Algorithms to replace divisions by a constant by multiplications and shifts are given in [134], see also [265].

We note that the C standard leaves the behavior of a right shift of a signed integer as 'implementation-defined'. The described behavior (that a negative value remains negative after right shift) is the de facto standard of all modern C compilers.

1.1.4 A pitfall (two's complement)

```
 c=................  -c=................   c=     0  -c=     0  <--=
 c=...............1  -c=1111111111111111   c=     1  -c=    -1
 c=..............1.  -c=111111111111111.   c=     2  -c=    -2
 c=..............11  -c=111111111111111.1  c=     3  -c=    -3
 c=.............1..  -c=11111111111111..   c=     4  -c=    -4
 c=.............1.1  -c=11111111111111.11  c=     5  -c=    -5
 c=.............11.  -c=11111111111111.1.  c=     6  -c=    -6
[--snip--]
 c=.111111111111.1  -c=1............11   c= 32765  -c=-32765
 c=.111111111111.1  -c=1............1.   c= 32766  -c=-32766
 c=.111111111111.1  -c=1............1    c= 32767  -c=-32767
 c=1................  -c=1...............   c=-32768  -c=-32768  <--=
 c=1...............1  -c=.111111111111111   c=-32767  -c= 32767
 c=1..............1.  -c=.111111111111111.  c=-32766  -c= 32766
 c=1..............11  -c=.111111111111111.1 c=-32765  -c= 32765
 c=1.............1..  -c=.11111111111111..  c=-32764  -c= 32764
 c=1.............1.1  -c=.11111111111111.11 c=-32763  -c= 32763
 c=1.............11.  -c=.11111111111111.1. c=-32762  -c= 32762
[--snip--]
 c=111111111111..1  -c=.............111   c=    -7  -c=     7
 c=111111111111.1.  -c=.............11.   c=    -6  -c=     6
 c=111111111111.11  -c=.............1.1   c=    -5  -c=     5
 c=111111111111...  -c=.............1..   c=    -4  -c=     4
 c=111111111111..1  -c=.............11    c=    -3  -c=     3
 c=111111111111.1.  -c=.............1.    c=    -2  -c=     2
 c=111111111111.11  -c=.............1     c=    -1  -c=     1
```

Figure 1.1-A: With two's complement there is one nonzero value that is its own negative.

In two's complement zero is not the only number that is equal to its negative. With a data type of n bits the value with just the highest bit set (the most negative value) also has this property. Figure 1.1-A (the output of [FXT: bits/gotcha-demo.cc]) shows the situation for words of sixteen bits. This is the reason why innocent looking code like

```
if ( x<0 )  x = -x;
// assume x positive here (WRONG!)
```

can simply fail.

1.1.5 Another pitfall (shifts in the C-language)

A shift by more than `BITS_PER_LONG`−1 is undefined by the C-standard. Therefore the following function can fail if `k` is zero:

```
1  inline ulong first_comb(ulong k)
2  // Return the first combination of (i.e. smallest word with) k bits,
3  // i.e.  00..001111..1 (k low bits set)
4  {
5      ulong t = ~0UL >> ( BITS_PER_LONG - k );
6      return t;
7  }
```

Compilers usually emit just a shift instruction which on certain CPUs does *not* give zero if the shift is equal to or greater than `BITS_PER_LONG`. This is why the line

```
if ( k==0 )  t = 0;  // shift with BITS_PER_LONG is undefined
```

has to be inserted just before the `return` statement.

1.1.6 Shortcuts

To test whether at least one of `a` and `b` equals zero use `if ( !(a && b) )`. This works for signed and unsigned integers. Checking whether both are zero can be done using `if ( (a|b)==0 )`. This obviously generalizes for several variables as `if ( (a|b|c|..|z)==0 )`. Test whether exactly one of two variables is zero using `if ( (!a) ^ (!b) )` .

1.1.7 Toggling between values

In order to toggle an integer `x` between two values `a` and `b` use:

```
precalculate:  t  = a ^ b;
toggle:        x ^= t;   // a <--> b
```

the equivalent trick for floating point types is

```
precalculate:  t = a + b;
toggle:        x = t - x;
```

Here an overflow could occur with `a` and `b` in the legal range (but both close to overflow). This should, however, not happen with sane programs.

1.1.8 Next or previous even or odd value

Compute the next or previous even or odd value via [FXT: bits/evenodd.h]:

```
1   static inline ulong next_even(ulong x)  { return x+2-(x&1); }
2   static inline ulong prev_even(ulong x)  { return x-2+(x&1); }
3
4   static inline ulong next_odd(ulong x)  { return x+1+(x&1); }
5   static inline ulong prev_odd(ulong x)  { return x-1-(x&1); }
```

The following functions return the unmodified argument if it has the required property, else the nearest such value:

```
1   static inline ulong next0_even(ulong x)  { return x+(x&1); }
2   static inline ulong prev0_even(ulong x)  { return x-(x&1); }
3
4   static inline ulong next0_odd(ulong x)  { return x+1-(x&1); }
5   static inline ulong prev0_odd(ulong x)  { return x-1+(x&1); }
```

1.1.9 Testing whether bit-subset

The following function tests whether a word u, as a bit-set, is a subset of another word e [FXT: bits/bitsubsetq.h]:

```
1   inline bool is_subset(ulong u, ulong e)
2   // Return whether u is a bit-subset of e.
3   {
4       return  ( (u & e)==u );
5   }
```

Should u contain any bits not set in e then these bits are deleted in the AND-operation and the test for equality will fail.

1.1.10 Integer versus float multiplication

The floating point multiplier gives the highest bits of the product. Integer multiplication gives the result modulo 2^b where b is the number of bits of the integer type used. As an example we square the number 1010101 using a 32-bit integer type and floating point types with 24-bit and 53-bit mantissa:

```
  a = 111111111
a*a = 12345678987654321 // true result

a*a = 1653732529            // result with 32-bit integer multiplication
(a*a)%(2**32) = 1653732529  // ... which is modulo (2**bits_per_int)

a*a = 1.2345679481405440e+16 // result with float multiplication (24 bit mantissa)
a*a = 1.2345678987654320e+16 // result with float multiplication (53 bit mantissa)
```

1.1.11 Double precision float to signed integer conversion

Conversion of double precision floats that have a 53-bit mantissa to signed integers via [13, p.52-53]

```
#define DOUBLE2INT(i, d)  { double t = ((d) + 6755399441055744.0); i = *((int *)(&t)); }
double x;
int i = 123;
DOUBLE2INT(i, x);
```

can be a faster alternative to

```
double x = 123.0;
int i;
i = x;
```

The constant used is $6755399441055744 = 2^{52} + 2^{51}$. The method is machine dependent as it relies on the binary representation of the floating point mantissa. Here it is assumed that, firstly, the floating point number has a 53-bit mantissa with the most significant bit (that is always one with normalized numbers) omitted, and secondly, the address of the number points to the mantissa.

1.1.12 Optimization considerations

Never ever think that some code is the 'fastest possible', there always another trick that can still improve performance. Many factors can have an influence on performance like number of CPU registers or cost of branches. Code that performs well on one machine might perform badly on another. The old trick to swap variables without using a temporary

```
          //   a=0, b=0    a=0, b=1    a=1, b=0    a=1, b=1
a ^= b; //     0    0      1    1      1    0      0    1
b ^= a; //     0    0      1    0      1    1      0    1
a ^= b; //     0    0      1    0      0    1      1    1

equivalent to:
tmp = a;  a = b;  b = tmp;
```

is pretty much out of fashion today. However in some specific context (like extreme register pressure) it may be the way to go.

The only way to find out which version of a function is faster is to actually do profiling (timing). The performance does depend on the stream of instructions before the machine code (we assume that all of these low-level functions get inlined). Studying the generated CPU instructions does help to understand what is going on but can never replace profiling.

The code surrounding a specific function can have a massive impact on performance. That is, benchmarks for just the isolated routine can only give a rough indication. Profile your application and also test whether the second best (when isolated) routine is the fastest.

Never just replace the unoptimized version of some code fragment when introducing a streamlined one. Keep the original in the source. In case something nasty happens (think of low level software failures when porting to a different platform) you'll be very thankful for the chance to temporarily use the slow but correct version.

Study the optimization recommendations for your CPU (like [13] for the AMD64). It doesn't hurt to see the corresponding documentation for other architectures.

Proper documentation is an absolute must for optimized code, just assume that nobody will be able to read and understand it from the supplied source alone. The experience of not being able to understand code you have written some time ago helps a lot in this matter.

More techniques for optimization are given in section 5.4 on page 172.

1.2 Operations on individual bits

1.2.1 Testing, setting, and deleting bits

The following functions should be self explanatory. Following the spirit of the C language there is no check whether the indices used are out of bounds. That is, if any index is greater or equal `BITS_PER_LONG`, the result is undefined [FXT: bits/bittest.h]:

```
1  inline ulong test_bit(ulong a, ulong i)
2  // Return zero if bit[i] is zero,
3  //  else return one-bit word with bit[i] set.
4  {
5      return  (a & (1UL << i));
6  }
```

The following version returns either zero or one:

```
1  static inline bool test_bit01(ulong a, ulong i)
2  // Return whether bit[i] is set.
3  {
4      return ( 0 != test_bit(a, i) );
5  }
```

```
1  inline ulong set_bit(ulong a, ulong i)
2  // Return a with bit[i] set.
3  {
4      return  (a | (1UL << i));
5  }
```

```
1  inline ulong delete_bit(ulong a, ulong i)
2  // Return a with bit[i] cleared.
3  {
4      return  (a & ~(1UL << i));
5  }
```

```
1  inline ulong change_bit(ulong a, ulong i)
2  // Return a with bit[i] changed.
3  {
4      return  (a ^ (1UL << i));
5  }
```

1.2.2 Copying a bit

In order to copy a bit from one position to another we generate a one exactly if the bits at the two positions differ. Then an XOR changes the target bit if needed [FXT: bits/bitcopy.h]:

```
1  inline ulong copy_bit(ulong a, ulong isrc, ulong idst)
2  // Copy bit at [isrc] to position [idst].
3  // Return the modified word.
4  {
5      ulong x = ((a>>isrc) ^ (a>>idst)) & 1;  // one if bits differ
6      a ^= (x<<idst); // change if bits differ
7  }
```

The situation is more tricky if the bit positions are given as (one bit) masks:

```
1  inline ulong mask_copy_bit(ulong a, ulong msrc, ulong mdst)
2  // Copy bit according at src-mask (msrc)
3  //  to the bit according to the dest-mask (mdst).
4  // Both msrc and mdst must have exactly one bit set.
5  // Return the modified word.
6  {
```

```
7       ulong x = mdst;
8       if ( msrc & a )  x = 0;  // zero if source bit set
9       x ^= mdst;  // ==mdst if source bit set, else zero
10      a &= ~mdst;  // clear dest bit
11      a |= x;
12      return a;
13  }
```

The compiler generates branch-free code as the conditional assignment is compiled to a `cmov` (conditional move) assembler instruction. If one or both masks have several bits set the routine will set all bits of `mdst` if any of the bits in `msrc` is one else clear all bits of `mdst`.

1.2.3 Swapping two bits

A function to swap two bits of a word [FXT: bits/bitswap.h]:

```
1   static inline ulong bit_swap(ulong a, ulong k1, ulong k2)
2   // Return a with bits at positions [k1] and [k2] swapped.
3   // k1==k2 is allowed (a is unchanged then)
4   {
5       ulong x = ((a>>k1) ^ (a>>k2)) & 1;  // one if bits differ
6       a ^= (x<<k2); // change if bits differ
7       a ^= (x<<k1); // change if bits differ
8       return  a;
9   }
```

When it is known that the bits do have different values the following routine can be used:

```
1   static inline ulong bit_swap_01(ulong a, ulong k1, ulong k2)
2   // Return a with bits at positions [k1] and [k2] swapped.
3   // Bits must have different values (!)
4   // (i.e. one is zero, the other one)
5   // k1==k2 is allowed (a is unchanged then)
6   {
7       return  a ^ ( (1UL<<k1) ^ (1UL<<k2) );
8   }
```

1.3 Operations on low bits or blocks of a word

The underlying idea of functions that operate on the lowest set bit is that addition and subtraction of 1 always changes a burst of bits at the lower end of the word. The following functions are given in [FXT: bits/bitlow.h].

Isolation of the lowest set bit is achieved via

```
1   static inline ulong lowest_one(ulong x)
2   // Return word where only the lowest set bit in x is set.
3   // Return 0 if no bit is set.
4   {
5       return  x & -x;  // use: -x == ~x + 1
6   }
```

The lowest zero (unset bit) of some word `x` is then trivially isolated using the equivalent of `lowest_one( ~x )`:

```
1   static inline ulong lowest_zero(ulong x)
2   // Return word where only the lowest unset bit in x is set.
3   // Return 0 if all bits are set.
4   {
5       x = ~x;
6       return  x & -x;
7   }
```

Alternatively, one can use either of

```
return  (x ^ (x+1)) & ~x;
return  ((x ^ (x+1)) >> 1 ) + 1;
```

The sequence of returned values for $x = 0, 1, \ldots$ is the binary ruler function, the highest power of two that divides $x + 1$:

```
   x:  ==   x        lowest_zero(x)
   0:  == ........     .......1
   1:  == .......1     ......1.
   2:  == ......1.     .......1
   3:  == ......11     .....1..
   4:  == .....1..     .......1
   5:  == .....1.1     ......1.
   6:  == .....11.     .......1
   7:  == .....111     ....1...
   8:  == ....1...     .......1
   9:  == ....1..1     ......1.
  10:  == ....1.1.     .......1
```

Clearing the lowest set bit in a word can be achieved via

```
1   static inline ulong delete_lowest_one(ulong x)
2   // Return word where the lowest bit set in x is cleared.
3   // Return 0 for input == 0.
4   {
5       return  x & (x-1);
6   }
```

while setting the lowest unset bit is done by

```
1   static inline ulong set_lowest_zero(ulong x)
2   // Return word where the lowest unset bit in x is set.
3   // Return ~0 for input == ~0.
4   {
5       return  x | (x+1);
6   }
```

Isolate the burst of low ones or zeros as follows:

```
1   static inline ulong low_ones(ulong x)
2   // Return word where all the (low end) ones are set.
3   // Example:  01011011 --> 00000011
4   // Return 0 if lowest bit is zero:
5   //       10110110 --> 0
6   {
7       x = ~x;
8       x &= -x;
9       --x;
10      return x;
11  }
```

and

```
1   static inline ulong low_zeros(ulong x)
2   // Return word where all the (low end) zeros are set.
3   // Example:  01011000 --> 00000111
4   // Return 0 if all bits are set.
5   {
6       x &= -x;
7       --x;
8       return x;
9   }
```

Isolation of the lowest block of ones (which may have zeros to the right of it) can be achieved via:

```
1   static inline ulong lowest_block(ulong x)
2   // Isolate lowest block of ones.
3   // e.g.:
4   // x   = *****011100
5   // l   = 00000000100
6   // y   = *****100000
7   // x^y = 00000111100
8   // ret = 00000011100
9   {
10      ulong l = x & -x;  // lowest bit
11      ulong y = x + l;
12      x ^= y;
13      return  x & (x>>1);
14  }
```

Extracting the *index* (position) of the lowest bit is easy when the corresponding assembler instruction is used [FXT: bits/bitasm-amd64.h]:

```
1   static inline ulong asm_bsf(ulong x)
```

```
2   // Bit Scan Forward
3   {
4       asm ("bsfq %0, %0" : "=r" (x) : "0" (x));
5       return x;
6   }
```

Without the assembler instruction an algorithm that uses proportional $\log_2$(BITS_PER_LONG) can be used, so the resulting function can be implemented as[2] (64-bit version)

```
1   static inline ulong lowest_one_idx(ulong x)
2   // Return index of lowest bit set.
3   // Examples:
4   //    ***1 --> 0
5   //    **10 --> 1
6   //    *100 --> 2
7   // Return 0 (also) if no bit is set.
8   {
9       ulong r = 0;
10      x &= -x;  // isolate lowest bit
11      if ( x & 0xffffffff00000000UL )  r += 32;
12      if ( x & 0xffff0000ffff0000UL )  r += 16;
13      if ( x & 0xff00ff00ff00ff00UL )  r += 8;
14      if ( x & 0xf0f0f0f0f0f0f0f0UL )  r += 4;
15      if ( x & 0xccccccccccccccccUL )  r += 2;
16      if ( x & 0xaaaaaaaaaaaaaaaaUL )  r += 1;
17
18  #endif // BITS_USE_ASM
19      return r;
20  }
```

The function returns zero for two inputs, one and zero. If one needs a special value for the input zero, add a statement like

```
if ( 1>=x )  return  x-1; // 0 if 1, ~0 if 0
```

as first line of the function.

Occasionally one wants to set a rising or falling edge at the position of the lowest bit:

```
1   static inline ulong lowest_one_01edge(ulong x)
2   // Return word where a all bits from  (including) the
3   //   lowest set bit to bit 0 are set.
4   // Return 0 if no bit is set.
5   {
6       if ( 0==x )  return 0;
7       return  x^(x-1);
8   }
```

```
1   static inline ulong lowest_one_10edge(ulong x)
2   // Return word where a all bits from (including) the
3   //   lowest set bit to most significant bit are set.
4   // Return 0 if no bit is set.
5   {
6       if ( 0==x )  return 0;
7       x ^= (x-1);
8       // here  x == lowest_one_01edge(x);
9       return  ~(x>>1);
10  }
```

The following function returns the parity of the lowest bit in a binary word

```
1   static inline ulong lowest_one_idx_parity(ulong x)
2   {
3       x &= -x;  // isolate lowest bit
4       return  (x & 0xaaaaaaaaaaaaaaaaUL);
5   }
```

The sequence of values for $x = 0, 1, 2, \ldots$ is

```
00100010101000100010001010100010101000101010001000100010101000010...
```

This is the complement of the *period-doubling sequence*, entry A035263 of [240]. See section 36.5.1 on page 723 for the connection to the towers of Hanoi puzzle.

[2] thanks go to Nathan Bullock for communicating this improved version.

1.4 Extraction of ones, zeros, or blocks near transitions

We give functions for the creation or extraction of bit-blocks, and the isolation of values near transitions. The term 'one' is used for a set bit, and 'zero' for an unset bit. A transition is a place where adjacent bits have different values. A block is a goup of at least two adjacent bits of the same value.

1.4.1 Creating blocks of ones

The following functions are given in [FXT: bits/bitblock.h].

```
1  static inline ulong bit_block(ulong p, ulong n)
2  // Return word with length-n bit block starting at bit p set.
3  // Both p and n are effectively taken modulo BITS_PER_LONG.
4  {
5      ulong x = (1UL<<n) - 1;
6      return  x << p;
7  }
```

A version with indices wrapping around is

```
1  static inline ulong cyclic_bit_block(ulong p, ulong n)
2  // Return word with length-n bit block starting at bit p set.
3  // The result is possibly wrapped around the word boundary.
4  // Both p and n are effectively taken modulo BITS_PER_LONG.
5  {
6      ulong x = (1UL<<n) - 1;
7      return  (x<<p) | (x>>(BITS_PER_LONG-p));
8  }
```

1.4.2 Finding isolated ones or zeros

For the following functions [FXT: bits/bitmisc.h]. we assume that the outside bits are all zero

```
1  static inline ulong single_ones(ulong x)
2  // Return word with only the isolated ones from x set.
3  {
4      return  x & ~( (x<<1) | (x>>1) );
5  }
```

```
1  static inline ulong single_zeros(ulong x)
2  // Return word with only the isolated zeros from x set.
3  {
4      return  single_ones( ~x );
5  }
```

```
1  static inline ulong single_values(ulong x)
2  // Return word where only the isolated ones and zeros from x are set.
3  {
4      return  (x ^ (x<<1)) & (x ^ (x>>1));
5  }
```

1.4.3 Isolating single ones or zeros at the word boundary

```
1  static inline ulong border_ones(ulong x)
2  // Return word where only those ones from x are set that lie next to a zero.
3  {
4      return  x & ~( (x<<1) & (x>>1) );
5  }
```

```
1  static inline ulong border_values(ulong x)
2  // Return word where those bits from x are set that lie on a transition.
3  {
4      ulong g = x ^ (x>>1);
5      g |= (g<<1);
6      return   g | (x & 1);
7  }
```

1.4.4 Isolating transitions

```
1  static inline ulong high_border_ones(ulong x)
2  // Return word where only those ones from x are set
3  //   that lie right to (i.e. in the next lower bin of) a zero.
4  {
5      return  x & ( x ^ (x>>1) );
6  }
```

```
1  static inline ulong low_border_ones(ulong x)
2  // Return word where only those ones from x are set
3  //   that lie left to (i.e. in the next higher bin of) a zero.
4  {
5      return  x & ( x ^ (x<<1) );
6  }
```

1.4.5 Isolating ones or zeros at block boundaries

```
1  static inline ulong block_border_ones(ulong x)
2  // Return word where only those ones from x are set
3  //   that are at the border of a block of at least 2 bits.
4  {
5      return  x & ( (x<<1) ^ (x>>1) );
6  }
```

```
1  static inline ulong low_block_border_ones(ulong x)
2  // Return word where only those bits from x are set
3  //   that are at left of a border of a block of at least 2 bits.
4  {
5      ulong t = x & ( (x<<1) ^ (x>>1) );  // block_border_ones()
6      return  t & (x>>1);
7  }
```

```
1  static inline ulong high_block_border_ones(ulong x)
2  // Return word where only those bits from x are set
3  //   that are at right of a border of a block of at least 2 bits.
4  {
5      ulong t = x & ( (x<<1) ^ (x>>1) );  // block_border_ones()
6      return  t & (x<<1);
7  }
```

```
1  static inline ulong block_ones(ulong x)
2  // Return word where only those bits from x are set
3  //   that are part of a block of at least 2 bits.
4  {
5      return  x & ( (x<<1) | (x>>1) );
6  }
```

1.4.6 Isolating the interior of bit blocks

```
1  static inline ulong block_values(ulong x)
2  // Return word where only those bits are set
3  // that do not lie next to an opposite value.
4  {
5      return  ~single_values(x);
6  }
```

```
1  static inline ulong interior_ones(ulong x)
2  // Return word where only those bits from x are set
3  //   that do not have a zero to their left or right.
4  {
5      return  x & ( (x<<1) & (x>>1) );
6  }
```

```
1  static inline ulong interior_values(ulong x)
2  // Return word where only those values from x are set
3  //   that do have a transitions on both sides.
4  {
5      return  ~border_values(x);
6  }
```

1.5 Computing the index of a single set bit

In the function `lowest_one_idx()` we first isolated the lowest bit of a word `x` by first setting `x&=-x`. At this point, `x` contains just one set bit (or `x==0`). The following lines in the routine implement an algorithm that computes the index of the single bit set. This section gives some alternative techniques to compute the index of a single-bit word.

1.5.1 Cohen's trick

A nice trick is presented in [90]: for N-bit words find a number m so that all powers of two are different modulo m. That is, the order of two modulo m must be greater or equal to N. We use a table `mt[]` of size m that contains the powers of two: `mt[(2**j) mod m] = j` for $j > 0$ and a special value for $j = 0$. To look up the index of a one-bit-word `x` it is reduced modulo m and `mt[x]` is returned.

```
modulus m=11
   k =  0  1  2  3  4  5  6  7
mt[k]=  0  0  1  8  2  4  9  7

Lowest bit == 0:       x= .......1 =   1    x % m=  1 ==> lookup = 0
Lowest bit == 1:       x= ......1. =   2    x % m=  2 ==> lookup = 1
Lowest bit == 2:       x= .....1.. =   4    x % m=  4 ==> lookup = 2
Lowest bit == 3:       x= ....1... =   8    x % m=  8 ==> lookup = 3
Lowest bit == 4:       x= ...1.... =  16    x % m=  5 ==> lookup = 4
Lowest bit == 5:       x= ..1..... =  32    x % m= 10 ==> lookup = 5
Lowest bit == 6:       x= .1...... =  64    x % m=  9 ==> lookup = 6
Lowest bit == 7:       x= 1....... = 128    x % m=  7 ==> lookup = 7
```

Figure 1.5-A: Determination of the position of a single bit with 8-bit words.

We demonstrate the method for $N = 8$ where $m = 11$ is the smallest number with the required property. The setup routine for the table is

```
1   const ulong m = 11; // the modulus
2   ulong mt[m+1];
3   static void mt_setup()
4   {
5       mt[0] = 0;  // special value for the zero word
6       ulong t = 1;
7       for (ulong i=1; i<m; ++i)
8       {
9           mt[t] = i-1;
10          t *= 2;
11          if ( t>=m )  t -= m;  // modular reduction
12      }
13  }
```

The entry in `mt[0]` will be accessed when the input is the zero word. One can use a special value that the algorithm will return for input zero. Here we simply used zero in order to always have the same return value as with `lowest_one_idx()`. The computation of the index can then be achieved by

```
1   inline ulong m_lowest_one_idx(ulong x)
2   {
3       x &= -x;  // isolate lowest bit
4       x %= m;   // power of two modulo m
5       return  mt[x];  // lookup
6   }
```

The code is given in the demo [FXT: bits/modular-lookup-demo.cc], the output with $N = 8$ (edited for size) is shown in figure 1.5-A. The following moduli $m(N)$ can be used for N-bit words:

```
N:  4   8  16   32   64   128   256    512    1024
m:  5  11  19   37   67   131   269    523    1061
```

The modulus $m(N)$ is the smallest prime greater than N such that 2 is a primitive root modulo $m(N)$:

```
for (n=2, 10,  N=2^n;   \\ N bits per word
    forprime (z=N, N+9999,
```

```
            if ( znorder(Mod(2,z))>=N, print(N,": ",z);break() )
        )
    )
```

1.5.2 Using De Bruijn sequences

The following method (given in [181]) is even more elegant, it uses binary De Bruijn sequences of size N. A binary De Bruijn sequence of length 2^N contains all binary words of length N (see section 39.1 on page 857). These are the sequences for 32 and 64 bit, as binary words:

```
#if BITS_PER_LONG == 32
const ulong db = 0x4653ADFUL;
// == 00000100011001010011101011011111
const ulong s = 32-5;
#else
const ulong db = 0x218A392CD3D5DBFUL;
// == 0000001000011000101000111001001011001101001111010101110110111111
const ulong s = 64-6;
#endif
```

```
db=...1.111  (De Bruijn sequence)
    k  =  0  1  2  3  4  5  6  7
dbt[k] =  0  1  2  4  7  3  6  5

Lowest bit == 0:    x = .......1   db * x = ...1.111   shifted = ........ == 0 ==> lookup = 0
Lowest bit == 1:    x = ......1.   db * x = ..1.111.   shifted = .......1 == 1 ==> lookup = 1
Lowest bit == 2:    x = .....1..   db * x = .1.111..   shifted = ......1. == 2 ==> lookup = 2
Lowest bit == 3:    x = ....1...   db * x = 1.111...   shifted = .....1.1 == 5 ==> lookup = 3
Lowest bit == 4:    x = ...1....   db * x = .111....   shifted = ......11 == 3 ==> lookup = 4
Lowest bit == 5:    x = ..1.....   db * x = 111.....   shifted = .....111 == 7 ==> lookup = 5
Lowest bit == 6:    x = .1......   db * x = 11......   shifted = .....11. == 6 ==> lookup = 6
Lowest bit == 7:    x = 1.......   db * x = 1.......   shifted = .....1.. == 4 ==> lookup = 7
```

Figure 1.5-B: Computing the position of the single set bit in 8-bit words with a De Bruijn sequence.

Let w_i be the i-th sub-word from the left (high end). We create a table so that the entry with index w_i points to i:

```
1   ulong dbt[BITS_PER_LONG];
2   static void dbt_setup()
3   {
4       for (ulong i=0; i<BITS_PER_LONG; ++i)  dbt[ (db<<i)>>s ] = i;
5   }
```

The computation of the index involves a multiplication and a table lookup:

```
1   inline ulong db_lowest_one_idx(ulong x)
2   {
3       x &= -x;  // isolate lowest bit
4       x *= db;  // multiplication by a power of two is a shift
5       x >>= s;  // use log_2(BITS_PER_LONG) highest bits
6       return  dbt[x];  // lookup
7   }
```

The used sequences must start with at least $\log_2(N) - 1$ zeros because in the line `x *= db` the word `x` is shifted (not rotated). The code is given in the demo [FXT: bits/debruijn-lookup-demo.cc], the output with $N = 8$ (edited for size, dots denote zeros) is shown in figure 1.5-B.

1.5.3 Using floating point numbers

Floating point numbers are normalized so that the highest bit in the mantissa is one. Therefore if one converts an integer into a float then the position of the *highest* set bit can be read off the exponent. By isolating the lowest bit before that operation its index can be found by the same trick. However, the conversion between integers and floats is usually slow. Further, the technique is highly machine dependent.

1.6 Operations on high bits or blocks of a word

For the functions operating on the highest bit there is not a way as trivial as with the equivalent task with the lower end of the word. With a bit-reverse CPU-instruction available life would be significantly easier. However, almost no CPU seems to have it. The following functions are given in [FXT: bits/bithigh.h].

Isolation of the highest set bit is achieved via the bit-scan instruction when it is available [FXT: bits/bitasm-i386.h]:

```
1   static inline ulong asm_bsr(ulong x)
2   // Bit Scan Reverse
3   {
4       asm ("bsrl %0, %0" : "=r" (x) : "0" (x));
5       return x;
6   }
```

else one may use

```
1   static inline ulong highest_one_01edge(ulong x)
2   // Return word where a all bits from (including) the
3   //   highest set bit to bit 0 are set.
4   // Return 0 if no bit is set.
5   {
6       x |= x>>1;
7       x |= x>>2;
8       x |= x>>4;
9       x |= x>>8;
10      x |= x>>16;
11  #if  BITS_PER_LONG >= 64
12      x |= x>>32;
13  #endif
14      return  x;
15  }
```

so the resulting code is

```
1   static inline ulong highest_one(ulong x)
2   // Return word where only the highest bit in x is set.
3   // Return 0 if no bit is set.
4   {
5   #if defined  BITS_USE_ASM
6       if ( 0==x )  return 0;
7       x = asm_bsr(x);
8       return  1UL<<x;
9   #else
10      x = highest_one_01edge(x);
11      return  x ^ (x>>1);
12  #endif // BITS_USE_ASM
13  }
```

Trivially, the highest zero can be isolated using `highest_one(~x)`. Thereby

```
1   static inline ulong set_highest_zero(ulong x)
2   // Return word where the highest unset bit in x is set.
3   // Return ~0 for input == ~0.
4   {
5       return  x | highest_one( ~x );
6   }
```

Finding the index of the highest set bit uses the equivalent algorithm as with the lowest set bit:

```
1   static inline ulong highest_one_idx(ulong x)
2   // Return index of highest bit set.
3   // Return 0 if no bit is set.
4   {
5   #if defined  BITS_USE_ASM
6       return  asm_bsr(x);
7   #else // BITS_USE_ASM
8
9       if ( 0==x )  return  0;
10
11      ulong r = 0;
12  #if  BITS_PER_LONG >= 64
13      if ( x & 0xffffffff00000000UL )  { x >>= 32;  r += 32; }
14  #endif
15      if ( x & 0xffff0000UL )  { x >>= 16;  r += 16; }
```

```
        ................1111....1111.111 = 0xf0f7  == word
        ................1............... = highest_one
        ................111111111111111 = highest_one_01edge
        1111111111111111............... = highest_one_10edge
                                      15 = highest_one_idx
        ............................... = low_zeros
        ............................111 = low_ones
        ............................1.. = lowest_one
        ...............................1 = lowest_one_01edge
        11111111111111111111111111111111 = lowest_one_10edge
                                       0 = lowest_one_idx
        ............................111 = lowest_block
        ................1111....1111.11. = delete_lowest_one
        ............................1... = lowest_zero
        ................1111....11111111 = set_lowest_zero
        ............................... = high_ones
        1111111111111111............... = high_zeros
        1............................... = highest_zero
        1...............1111....1111.111 = set_highest_zero
-------------------------------------------------------------------------------
        1111111111111111....1111....1... = 0xffff0f08  == word
        1............................... = highest_one
        11111111111111111111111111111111 = highest_one_01edge
        1............................... = highest_one_10edge
                                      31 = highest_one_idx
        ............................111 = low_zeros
        ............................... = low_ones
        ............................1... = lowest_one
        ............................1111 = lowest_one_01edge
        11111111111111111111111111111...  = lowest_one_10edge
                                       3 = lowest_one_idx
        ............................1... = lowest_block
        1111111111111111....1111........ = delete_lowest_one
        ............................... 1 = lowest_zero
        1111111111111111....1111....1..1 = set_lowest_zero
        1111111111111111................ = high_ones
        ............................... = high_zeros
        ................1............... = highest_zero
        1111111111111111...1111....1... = set_highest_zero
```

Figure 1.6-A: Operations on the highest and lowest bits (and blocks) of a binary word for two different 32-bit input words. Dots denote zeros.

```
16      if ( x & 0x0000ff00UL )  { x >>=  8;  r +=  8; }
17      if ( x & 0x000000f0UL )  { x >>=  4;  r +=  4; }
18      if ( x & 0x0000000cUL )  { x >>=  2;  r +=  2; }
19      if ( x & 0x00000002UL )  {            r +=  1; }
20      return r;
21  #endif // BITS_USE_ASM
22  }
```

Isolation of the high zeros goes like

```
1   static inline ulong high_zeros(ulong x)
2   // Return word where all the (high end) zeros are set.
3   // e.g.:  00011001 --> 11100000
4   // Returns 0 if highest bit is set:
5   //        11011001 --> 00000000
6   {
7       x |= x>>1;
8       x |= x>>2;
9       x |= x>>4;
10      x |= x>>8;
11      x |= x>>16;
12  #if  BITS_PER_LONG >= 64
13      x |= x>>32;
14  #endif
15      return  ~x;
16  }
```

The high bits can be isolated using arithmetical right shift

```
1   static inline ulong high_ones(ulong x)
```

```
2   // Return word where all the (high end) ones are set.
3   // e.g.  11001011 --> 11000000
4   // Returns 0 if highest bit is zero:
5   //       01110110 --> 00000000
6   {
7       long y = (long)x;
8       y &= y>>1;
9       y &= y>>2;
10      y &= y>>4;
11      y &= y>>8;
12      y &= y>>16;
13  #if  BITS_PER_LONG >= 64
14      y &= y>>32;
15  #endif
16      return  (ulong)y;
17  }
```

In case arithmetical shifts are more expensive than unsigned shifts, instead use

```
1   static inline ulong high_ones(ulong x)
2   {
3       return  high_zeros( ~x );
4   }
```

A demonstration of selected functions operating on the highest or lowest bit (or block) of binary words is given in [FXT: bits/bithilo-demo.cc]. A part of the output is shown in figure 1.6-A.

1.7 Functions related to the base-2 logarithm

The following functions are given in [FXT: bits/bit2pow.h]. A the function that returns $\lfloor \log_2(x) \rfloor$ can be implemented using the obvious algorithm:

```
1   static inline ulong ld(ulong x)
2   // Return k so that 2^k <= x < 2^(k+1)
3   // If x==0 then 0 is returned (!)
4   {
5       ulong k = 0;
6       while ( x>>=1 )  { ++k; }
7       return k;
8   }
```

But this is the same as `highest_one_idx()`, so one can use

```
1   static inline ulong ld(ulong x)
2   {
3       return  highest_one_idx(x);
4   }
```

The bit-wise algorithm can be faster if the average result is known to be small.

The function `one_bit_q()` can be used to determine whether its argument is a power of two:

```
1   static inline bool one_bit_q(ulong x)
2   // Return whether x \in {1,2,4,8,16,...}
3   {
4       ulong m = x-1;
5       return  (((x^m)>>1) == m);
6   }
```

The following function does the same except that it returns `true` also for the zero argument:

```
1   static inline bool is_pow_of_2(ulong x)
2   // Return whether x == 0(!) or x == 2**k
3   {
4       return  !(x & (x-1));
5   }
```

Occasionally useful in FFT based computations (where the length of the available FFTs is often restricted to powers of two) are

```
1   static inline ulong next_pow_of_2(ulong x)
2   // Return x if x=2**k
```

```
3   // else return 2**ceil(log_2(x))
4   // Exception: returns 0 for x==0
5   {
6       if ( is_pow_of_2(x) )  return x;
7       x |= x >> 1;
8       x |= x >> 2;
9       x |= x >> 4;
10      x |= x >> 8;
11      x |= x >> 16;
12  #if BITS_PER_LONG == 64
13      x |= x >> 32;
14  #endif
15      return  x + 1;
16  }
```

and

```
1   static inline ulong next_exp_of_2(ulong x)
2   // Return k if x=2**k else return k+1.
3   // Exception: returns 0 for x==0.
4   {
5       if ( x <= 1 )  return 0;
6       return ld(x-1) + 1;
7   }
```

The following version may be faster than the one given above if assembler inlines are used:

```
1   static inline ulong next_pow_of_2(ulong x)
2   {
3       if ( is_pow_of_2(x) )  return x;
4       ulong n = 1UL<<ld(x);  // n<x
5       return  n<<1;
6   }
```

1.8 Counting bits and blocks of a word

If your CPU does not have a bit count instruction (sometimes called 'population count') then you might use an algorithm given in [FXT: bits/bitcount.h]. The following functions need proportional to $\log_2$(BITS_PER_LONG) operations:

```
1   static inline ulong bit_count(ulong x)
2   // Return number of bits set
3   {
4       x = (0x55555555UL & x) + (0x55555555UL & (x>> 1));  // 0-2 in 2 bits
5       x = (0x33333333UL & x) + (0x33333333UL & (x>> 2));  // 0-4 in 4 bits
6       x = (0x0f0f0f0fUL & x) + (0x0f0f0f0fUL & (x>> 4));  // 0-8 in 8 bits
7       x = (0x00ff00ffUL & x) + (0x00ff00ffUL & (x>> 8));  // 0-16 in 16 bits
8       x = (0x0000ffffUL & x) + (0x0000ffffUL & (x>>16));  // 0-31 in 32 bits
9       return x;
10  }
```

The underlying idea is to do a search via bit masks. The code can be improved to either

```
1       x  = ((x>>1) & 0x55555555UL) + (x & 0x55555555UL);  // 0-2 in 2 bits
2       x  = ((x>>2) & 0x33333333UL) + (x & 0x33333333UL);  // 0-4 in 4 bits
3       x  = ((x>>4) + x) & 0x0f0f0f0fUL;                   // 0-8 in 4 bits
4       x +=  x>> 8;                                        // 0-16 in 8 bits
5       x +=  x>>16;                                        // 0-32 in 8 bits
6       return  x & 0xff;
```

or (taken from [11])

```
1       x -=  (x>>1) & 0x55555555UL;
2       x  = ((x>>2) & 0x33333333UL) + (x & 0x33333333UL);
3       x  = ((x>>4) + x) & 0x0f0f0f0fUL;
4       x *= 0x01010101UL;
5       return  x>>24;
```

Which of the latter two versions is faster mainly depends on the speed of integer multiplication.

For 64-bit words the masks have to be adapted and one more step must be added (example corresponding to the second variant above):

```
1       x = ((x>>1) & 0x5555555555555555UL) + (x & 0x5555555555555555)UL;  // 0-2 in 2 bits
2       x = ((x>>2) & 0x3333333333333333UL) + (x & 0x3333333333333333)UL;  // 0-4 in 4 bits
3       x = ((x>>4) + x) & 0x0f0f0f0f0f0f0f0fUL;                           // 0-8 in 4 bits
4       x +=  x>> 8;                                                      // 0-16 in 8 bits
5       x +=  x>>16;                                                      // 0-32 in 8 bits
6       x +=  x>>32;                                                      // 0-64 in 8 bits
7       return  x & 0xff;
```

The following code (communicated by Johan Rönnblom [priv.comm.]) may be advantageous for 32-bit systems where loading constants is somewhat expensive:

```
1   inline uint CountBits32(uint a)
2   {
3       uint mask = 011111111111UL;  // Note: octal constant
4       a = (a - ((a&~mask)>>1)) - ((a>>2)&mask);
5       a += a>>3;
6       a = (a & 070707) + ((a>>18) & 070707);
7       a *= 010101;
8       return  ((a>>12) & 0x3f);
9   }
```

The following branchless algorithm can be useful when branches are expensive:

```
1   static inline ulong bit_count_01(ulong x)
2   // Return number of bits in a word
3   // for words of the special form 00...0001...11
4   {
5       ulong ct = 0;
6       ulong a;
7   #if  BITS_PER_LONG == 64
8       a = (x & (1UL<<32)) >> (32-5);  // test bit 32
9       x >>= a;  ct += a;
10  #endif
11      a = (x & (1UL<<16)) >> (16-4);  // test bit 16
12      x >>= a;  ct += a;
13
14      a = (x & (1UL<<8)) >> (8-3);  // test bit 8
15      x >>= a;  ct += a;
16
17      a = (x & (1UL<<4)) >> (4-2);  // test bit 4
18      x >>= a;  ct += a;
19
20      a = (x & (1UL<<2)) >> (2-1);  // test bit 2
21      x >>= a;  ct += a;
22
23      a = (x & (1UL<<1)) >> (1-0);  // test bit 1
24      x >>= a;  ct += a;
25
26      ct += x & 1; // test bit 0
27
28      return ct;
29  }
```

With a table of 256 bytes, holding the bit-counts of the numbers $0 \ldots 255$, the bits can be counted as follows:

```
1   ulong bit_count(ulong x)
2   {
3       unsigned char ct = 0;
4       ct += tab[ x & 0xff ];   x >>= 8;
5       ct += tab[ x & 0xff ];   x >>= 8;
6     [--snip--]  /* BYTES_PER_LONG times */
7       ct += tab[ x & 0xff ];
8       return ct;
9   }
```

However, while table driven methods tend to excel in synthetic banchmarks they can be very slow if they cause cache misses.

1.8.1 Sparse counting

When the (average input) word is known to have only a few bits set the following sparse count variant can be advantageous:

```
1   static inline ulong bit_count_sparse(ulong x)
```

```
2   // Return number of bits set.
3   {
4       ulong n = 0;
5       while ( x )  { ++n;  x &= (x-1); }
6       return  n;
7   }
```

The loop will execute once for each set bit. Partial unrolling of the loop should be an improvement for most cases:

```
1       ulong n = 0;
2       do
3       {
4           n += (x!=0);  x &= (x-1);
5           n += (x!=0);  x &= (x-1);
6           n += (x!=0);  x &= (x-1);
7           n += (x!=0);  x &= (x-1);
8       }
9       while ( x );
10      return  n;
```

If the number of bits is close to the maximal value then use the given routine with the complement:

```
1   static inline ulong bit_count_dense(ulong x)
2   // Return number of bits set.
3   // The loop (of bit_count_sparse()) will execute once for
4   //   each unset bit (i.e. zero) of x.
5   {
6       return BITS_PER_LONG - bit_count_sparse( ~x );
7   }
```

1.8.2 Counting blocks

The number of bit-blocks in a binary word can computed by the following function:

```
1   static inline ulong bit_block_count(ulong x)
2   // Return number of bit blocks.
3   // E.g.:
4   // ..1..11111...111.  -> 3
5   // ...1..11111...111  -> 3
6   // ......1.....1.1..  -> 3
7   // .........111.1111  -> 2
8   {
9       return  bit_count( (x^(x>>1)) ) / 2 + (x & 1);
10  }
```

Similarly, the number of blocks with two or more bits can be counted via:

```
1   static inline ulong bit_block_ge2_count(ulong x)
2   // Return number of bit blocks with at least 2 bits.
3   // E.g.:
4   // ..1..11111...111.  -> 2
5   // ...1..11111...111  -> 2
6   // ......1.....1.1..  -> 0
7   // .........111.1111  -> 2
8   {
9       return  bit_block_count( x & ( (x<<1) & (x>>1) ) );
10  }
```

1.8.3 GCC builtins *

Newer versions of the C compiler of the GNU Compiler Collection (GCC [119], starting with version 3.4) offer a function `__builtin_popcountl(ulong)` that counts the bits of an unsigned long integer.

We list a few such functions, taken from [120]:

```
int __builtin_ffs (unsigned int x)
    Returns one plus the index of the least significant 1-bit of x,
    or if x is zero, returns zero.

int __builtin_clz (unsigned int x)
    Returns the number of leading 0-bits in x, starting at the
    most significant bit position.  If x is 0, the result is undefined.
```

```
int __builtin_ctz (unsigned int x)
    Returns the number of trailing 0-bits in x, starting at the
    least significant bit position.  If x is 0, the result is undefined.
int __builtin_popcount (unsigned int x)
    Returns the number of 1-bits in x.
int __builtin_parity (unsigned int x)
    Returns the parity of x, i.e. the number of 1-bits in x modulo 2.
```

The names of corresponding versions for arguments of type unsigned long are obtained by adding '`l`' (ell) to the names.

1.9 Counting bits of many words

```
    x[ 0]=11111111  a0=11111111  a1=........  a2=........  a3=........  a4=........
    x[ 1]=11111111  a0=........  a1=11111111  a2=........  a3=........  a4=........
    x[ 2]=11111111  a0=11111111  a1=11111111  a2=........  a3=........  a4=........
    x[ 3]=11111111  a0=........  a1=........  a2=11111111  a3=........  a4=........
    x[ 4]=11111111  a0=11111111  a1=........  a2=11111111  a3=........  a4=........
    x[ 5]=11111111  a0=........  a1=11111111  a2=11111111  a3=........  a4=........
    x[ 6]=11111111  a0=11111111  a1=11111111  a2=11111111  a3=........  a4=........
    x[ 7]=11111111  a0=........  a1=........  a2=........  a3=11111111  a4=........
    x[ 8]=11111111  a0=11111111  a1=........  a2=........  a3=11111111  a4=........
    x[ 9]=11111111  a0=........  a1=11111111  a2=........  a3=11111111  a4=........
    x[10]=11111111  a0=11111111  a1=11111111  a2=........  a3=11111111  a4=........
    x[11]=11111111  a0=........  a1=........  a2=11111111  a3=11111111  a4=........
    x[12]=11111111  a0=11111111  a1=........  a2=11111111  a3=11111111  a4=........
    x[13]=11111111  a0=........  a1=11111111  a2=11111111  a3=11111111  a4=........
    x[14]=11111111  a0=11111111  a1=11111111  a2=11111111  a3=11111111  a4=........
    x[15]=11111111  a0=........  a1=........  a2=........  a3=........  a4=11111111
    x[16]=11111111  a0=11111111  a1=........  a2=........  a3=........  a4=11111111
```

Figure 1.9-A: Counting the bits of an array (where all bits are set) via vertical addition.

For counting the bits in a long array the technique of *vertical addition* can be useful. For ordinary addition the following relation holds:

```
a + b  ==  (a^b) + ((a&b)<<1)
```

The carry term `(a&b)` is propagated to the left. We can replace this 'horizontal' propagation by a 'vertical' one, that is, propagation into another word. The idea can be implemented as follows [FXT: bits/bitcount-v-demo.cc]:

```
1   ulong
2   bit_count_leq31(const ulong *x, ulong n)
3   // Return sum(j=0, n-1, bit_count(x[j]) )
4   // Must have  n<=31
5   {
6       ulong a0=0, a1=0, a2=0, a3=0, a4=0;
7       //       1,    3,    7,   15,   31,  <--= max n
8       for (ulong k=0; k<n; ++k)
9       {
10          ulong cy = x[k];
11          { ulong t = a0 & cy;  a0 ^= cy;  cy = t; }
12          { ulong t = a1 & cy;  a1 ^= cy;  cy = t; }
13          { ulong t = a2 & cy;  a2 ^= cy;  cy = t; }
14          { ulong t = a3 & cy;  a3 ^= cy;  cy = t; }
15          { a4 ^= cy; }
16          // [ PRINT x[k], a0, a1, a2, a3, a4 ]
17      }
18
19      ulong b = bit_count(a0);
20      b += (bit_count(a1)<<1);
21      b += (bit_count(a2)<<2);
22      b += (bit_count(a3)<<3);
23      b += (bit_count(a4)<<4);
24      return  b;
25  }
```

Figure 1.9-A shows the intermediate values with the computation of a length-17 array of all-ones words. After the loop the values of the variables `a0`, ..., `a4` are

```
a4=11111111
a3=........
a2=........
a1=
a0=iiiiiiii
```

Reading the columns as binary numbers tells us that in all positions of all words there were a total of $17 = 10001_2$ bits. The remaining instructions compute the total bit-count.

After some simplifications and loop-unrolling a routine for counting the bits of 15 words can be given as [FXT: bits/bitcount-v.cc]:

```
1   static inline ulong bit_count_v15(const ulong *x)
2   // Return sum(j=0, 14, bit_count(x[j]) )
3   // Technique is "vertical" addition.
4   {
5   #define VV(A) { ulong t = A & cy;  A ^= cy;  cy = t; }
6       ulong a1, a2, a3;
7       ulong a0=x[0];
8       { ulong cy = x[ 1];  VV(a0);  a1 = cy; }
9       { ulong cy = x[ 2];  VV(a0);  a1 ^= cy; }
10      { ulong cy = x[ 3];  VV(a0);  VV(a1);  a2 = cy; }
11      { ulong cy = x[ 4];  VV(a0);  VV(a1);  a2 ^= cy; }
12      { ulong cy = x[ 5];  VV(a0);  VV(a1);  a2 ^= cy; }
13      { ulong cy = x[ 6];  VV(a0);  VV(a1);  a2 ^= cy; }
14      { ulong cy = x[ 7];  VV(a0);  VV(a1);  VV(a2);  a3 = cy; }
15      { ulong cy = x[ 8];  VV(a0);  VV(a1);  VV(a2);  a3 ^= cy; }
16      { ulong cy = x[ 9];  VV(a0);  VV(a1);  VV(a2);  a3 ^= cy; }
17      { ulong cy = x[10];  VV(a0);  VV(a1);  VV(a2);  a3 ^= cy; }
18      { ulong cy = x[11];  VV(a0);  VV(a1);  VV(a2);  a3 ^= cy; }
19      { ulong cy = x[12];  VV(a0);  VV(a1);  VV(a2);  a3 ^= cy; }
20      { ulong cy = x[13];  VV(a0);  VV(a1);  VV(a2);  a3 ^= cy; }
21      { ulong cy = x[14];  VV(a0);  VV(a1);  VV(a2);  a3 ^= cy; }
22  #undef VV
23
24      ulong b = bit_count(a0);
25      b += (bit_count(a1)<<1);
26      b += (bit_count(a2)<<2);
27      b += (bit_count(a3)<<3);
28      return  b;
29  }
```

Each of the macros `VV` gives three machine instructions, namely AND, XOR, and one MOVE (assignment). The routine for the user is

```
1   ulong
2   bit_count_v(const ulong *x, ulong n)
3   // Return sum(j=0, n-1, bit_count(x[j]) )
4   {
5       ulong b = 0;
6       const ulong *xe = x + n + 1;
7       while ( x+15 < xe )  // process blocks of 15 elements
8       {
9           b += bit_count_v15(x);
10          x += 15;
11      }
12
13      // process remaining elements:
14      const ulong r = (ulong)(xe-x-1);
15      for (ulong k=0; k<r; ++k)  b+=bit_count(x[k]);
16
17      return  b;
18  }
```

Compared to the obvious method of bit-counting

```
1   ulong bit_count_v2(const ulong *x, ulong n)
2   {
3       ulong b = 0;
4       for (ulong k=0; k<n; ++k)  b += bit_count(x[k]);
5       return  b;
6   }
```

our routine uses roughly 30 percent less time when an array of 100,000,000 words is processed. There are many possible modifications of the method. If the bit-count routine is rather slow one may want to avoid the four calls to it after the processing of every 15 words. Instead, the variables `a0`, ..., `a3` could be added (vertically!) to an array of more elements. If that array has n elements then only with each block of $2^n - 1$ words n calls to the bit-count routine are necessary.

1.10 Bit set lookup

There is a nice trick to determine whether a given number is contained in a given subset of the set $\{0, 1, 2, \ldots, \texttt{BITS_PER_LONG}-1\}$. As an example, in order to determine whether `x` is a prime less than 32, one can use the function

```
ulong m = (1UL<<2) | (1UL<<3) | (1UL<<5) | ... | (1UL<<31);  // precomputed
static inline ulong is_tiny_prime(ulong x)
{
    return  m & (1UL << x);
}
```

The same idea applied to look up tiny factors [FXT: bits/tinyfactors.h]:

```
static inline bool is_tiny_factor(ulong x, ulong d)
// For x,d < BITS_PER_LONG (!)
// return whether d divides x  (1 and x included as divisors)
// no need to check whether d==0
//
{
    return  ( 0 != ( (tiny_factors_tab[x]>>d) & 1 ) );
}
```

The function uses the precomputed array [FXT: bits/tinyfactors.cc]:

```
extern const ulong tiny_factors_tab[] =
{
                    0x0UL,  // x = 0:           ( bits: ........)
                    0x2UL,  // x = 1:  1        ( bits: ......1.)
                    0x6UL,  // x = 2:  1 2      ( bits: .....11.)
                    0xaUL,  // x = 3:  1 3      ( bits: ....1.1.)
                   0x16UL,  // x = 4:  1 2 4    ( bits: ...1.11.)
                   0x22UL,  // x = 5:  1 5      ( bits: ..1...1.)
                   0x4eUL,  // x = 6:  1 2 3 6  ( bits: .1..111.)
                   0x82UL,  // x = 7:  1 7      ( bits: 1.....1.)
                  0x116UL,  // x = 8:  1 2 4 8
                  0x20aUL,  // x = 9:  1 3 9
[--snip--]
             0x20000002UL,  // x = 29:  1 29
             0x4000846eUL,  // x = 30:  1 2 3 5 6 10 15 30
             0x80000002UL,  // x = 31:  1 31
#if  ( BITS_PER_LONG > 32 )
            0x100010116UL,  // x = 32:  1 2 4 8 16 32
            0x20000080aUL,  // x = 33:  1 3 11 33
[--snip--]
    0x2000000000000002UL,  // x = 61:  1 61
    0x4000000080000006UL,  // x = 62:  1 2 31 62
    0x800000000020028aUL   // x = 63:  1 3 7 9 21 63
#endif // ( BITS_PER_LONG > 32 )
};
```

Bit arrays of arbitrary size are discussed in section 4.6 on page 160.

1.11 Avoiding branches

Branches are expensive operations with many CPUs, especially if the CPU pipeline is very long. A trick that is often useful is to replace

```
if ( (x<0) || (x>m) )  { ... }
```

where `x` might be a signed integer, by

```
if ( (unsigned)x > m )  { ... }
```

To test whether a point (x, y) lies outside square box whose size `m` is a power of two one should, instead of the sequence

```
if ( (x<0) || (x>m) || (y<0) || (y>m) )  { ... }
```

better use

```
if ( (x|y) > m )  { ... }
```

The following functions are given in [FXT: bits/branchless.h]. This function returns $\max(0, x)$. That is, zero is returned for negative input, else the unmodified input:

```
1   static inline long max0(long x)
2   {
3       return  x & ~(x >> (BITS_PER_LONG-1));
4   }
```

There is no restriction on input range. The trick used is that with negative x the arithmetic shift will give a word of all ones which is then negated and the AND-operation deletes all bits. Similarly:

```
1   static inline long min0(long x)
2   // Return min(0, x), i.e. return zero for positive input
3   {
4       return  x & (x >> (BITS_PER_LONG-1));
5   }
```

Computation of the average $(x + y)/2$ of two arguments x and y. The function gives the correct value even if $(x + y)$ does not fit into a machine word:

```
1   static inline ulong average(ulong x, ulong y)
2   // Return (x+y)/2
3   // Use the fact that x+y == ((x&y)<<1) + (x^y)
4   //    that is:       sum ==  carries   + sum_without_carries
5   {
6       return  (x & y) + ((x ^ y) >> 1);
7   }
```

If it is known that $x \geq y$ then one can alternatively use the statement `return y+(x-y)/2`.

The following `upos_*()` functions only work for a limited range. The highest bit must not be set in order to have the highest bit emulate the carry flag. Branchless computation of the absolute difference $|a - b|$:

```
1   static inline ulong upos_abs_diff(ulong a, ulong b)
2   {
3       long d1 = b - a;
4       long d2 = (d1 & (d1>>(BITS_PER_LONG-1)))<<1;
5       return  d1 - d2; // == (b - d) - (a + d);
6   }
```

Sorting of the arguments:

```
1   static inline void upos_sort2(ulong &a, ulong &b)
2   // Set {a, b} := {min(a, b), max(a,b)}
3   // Both a and b must not have the most significant bit set
4   {
5       long d = b - a;
6       d &= (d>>(BITS_PER_LONG-1));
7       a += d;
8       b -= d;
9   }
```

The following two functions adjust a given values when it lies outside a given range.

```
1   static inline long clip_range0(long x, long m)
2   // Code equivalent (for m>0) to:
3   //   if ( x<0 )  x = 0;
4   //   else if ( x>m )  x = m;
5   //   return  x;
6   {
7       if ( (ulong)x > (ulong)m )  x = m & ~(x >> (BITS_PER_LONG-1));
8       return  x;
9   }
```

```
1   static inline long clip_range(long x, long mi, long ma)
2   // Code equivalent to (for mi<=ma):
3   //   if ( x<mi )  x = mi;
4   //   else if ( x>ma )  x = ma;
5   {
6       x -= mi;
7       x = clip_range0(x, ma-mi);
8       x += mi;
9       return  x;
10  }
```

Johan Rönnblom gives the following versions for signed integer minimum, maximum, and absolute value, that can be advantageous for PPC (G4) CPUs:

```
1   #define B1  (BITS_PER_LONG-1) // bits of signed int minus one
2   #define MINI(x,y) (((x) &  (((int)((x)-(y)))>>B1)) + ((y) & ~(((int)((x)-(y)))>>B1)))
3   #define MAXI(x,y) (((x) & ~(((int)((x)-(y)))>>B1)) + ((y) &  (((int)((x)-(y)))>>B1))))
4   #define ABSI(x)   (((x) & ~(((int)(x))>>B1))       - ((x) &  (((int)(x))>>B1)))
```

1.11.1 Conditional swap

The following statement is compiled with a branch:

```
    if ( a<b )   { ulong t=a; a=b; b=t; } // swap if a < b
// Here:   a in  %rcx,    b in  %rdx
 62 000e 4889C8                movq    %rcx, %rax      # X, X
 68 0011 4839D1                cmpq    %rdx, %rcx      # X, X
 69 0014 7306                  jae     .L3     #,   // the branch
 71 0016 4889D1                movq    %rdx, %rcx      # X, X
 72 0019 4889C2                movq    %rax, %rdx      # X, X
 73                    .L3:
```

As conditional assignments can be done branchless, an equivalent branchless version is:

```
    { ulong x=a^b;  if (a>=b) { x=0; }  a^=x;  b^=x; } // swap if a < b
255 00af 4889EA                movq    %rbp, %rdx      # a, x
257 00b2 B9000000              movl    $0, %ecx        #, tmp83
257      00
260 00b7 4831C2                xorq    %rax, %rdx      # b, x
261 00ba 4839C5                cmpq    %rax, %rbp      # b, a
262 00bd 480F43D1              cmovae  %rcx, %rdx      # x,, tmp83, x
263 00c1 4831D5                xorq    %rdx, %rbp      # x, a
264 00c4 4831D0                xorq    %rdx, %rax      # x, b
```

We'd like to have fewer instructions. If one tries

```
   { ulong ta=a;  if (a<b) {a=b; b=ta;} } // swap if a < b
```

the generated code is identical to the first version. Let's try [FXT: bits/cswap.h]:

```
1   static inline void cswap_lt(ulong &a, ulong &b)
2   // Branchless equivalent to:
3   // if ( a<b )   { ulong t=a; a=b; b=t; } // swap if a < b
4   {
5       asm volatile("movq %0, %%r15 \n"   // t=a
6                    "cmpq %0, %1 \n"      // cmp a, b
7                    "cmovae %1, %0 \n"    // cond a=b
8                    "cmovae %%r15, %1 \n" // cond b=t
9                    ""
10                   : "=r" (a), "=r" (b) // output
11                   : "0" (a), "1" (b)   // input
12                   : "r15" // clobber
13                   );
14  }
```

Now the machine code looks better:

```
 // Here:   a in  %rax,    b in  %rdx
111 0027 4989C7        movq %rax, %r15         # tmp71
112 002a 4839C2        cmpq %rax, %rdx         # tmp71, tmp72
113 002d 480F43C2      cmovae %rdx, %rax       # tmp72, tmp71
114 0031 490F43D7      cmovae %r15, %rdx       # tmp72
```

Clearly, the relative speed of the three versions depends on the machine used. But it also turns out to be dependent on the surrounding code. We use bubble sort for benchmarking:

```
1  void bubble_sort(ulong *f, ulong n)
2  {
3      while ( n-- > 1 )
4          for (ulong k=0; k<n; ++k)  cswap_NN(f[k], f[k+1]);
5  }
```

Where we use the three versions of conditional swap for `cswap_NN()`. We sort an array of length 2^{15} twice with each version, once starting with an already sorted array and once with an array that is sorted in descending order. The 'plain' version wins:

```
cswap_gt_plain(f[k], f[k+1]); //  3.58s
cswap_gt_xor(f[k], f[k+1]);   //  6.34s
cswap_gt(f[k], f[k+1]);       //  5.10s
```

This is due to the fact that the compiler bypasses the store when no swap happens:

```
103 0020 488B4808  movq 8(%rax), %rcx #, tmp71
104 0024 488B10    movq (%rax), %rdx #* f, tmp72
105 0027 4839D1    cmpq %rdx, %rcx # tmp72, tmp71
106 002a 7307      jae .L7 #,
108 002c 48895008  movq %rdx, 8(%rax) # tmp72,
109 0030 488908    movq %rcx, (%rax) # tmp71,* f
111               .L7:
```

If we change the inner loop to

```
for (ulong k=0; k<n; ++k)
{
    ulong a = f[k],  b = f[k+1];
    cswap_NN(a, b);
    f[k] = a;   f[k+1] = b;
}
```

then we obtain:

```
cswap_gt_plain(a, b);  //  5.78s
cswap_gt_xor(a, b);    //  6.22s
cswap_gt(a, b);        //  4.68s
```

Our innocent looking change in the code prevented the compiler from doing its nice trick. The XOR version is (within timing precision) as slow as before. The assembler version wins because the data is already in registers. We learn that profiling is an absolute must.

1.11.2 Your compiler may be smarter than you thought

The machine code generated for

```
x =  x & ~(x >> (BITS_PER_LONG-1));  // max0()
```

is

```
 35:   48 99                 cqto
 37:   48 83 c4 08           add    $0x8,%rsp  // stack adjustment
 3b:   48 f7 d2              not    %rdx
 3e:   48 21 d0              and    %rdx,%rax
```

The variable `x` resides in the register `rAX` both at start and end of the function. The compiler uses a special (AMD64) instruction `cqto`. Quoting [12]:

> Copies the sign bit in the rAX register to all bits of the rDX register. The effect of this instruction is to convert a signed word, doubleword, or quadword in the rAX register into a signed doubleword, quadword, or double-quadword in the rDX:rAX registers. This action helps avoid overflow problems in signed number arithmetic.

Now the equivalent

```
x = ( x<0 ? 0 : x );  // max0()  "simple minded"
```

is compiled to:

```
  35:   ba 00 00 00 00          mov    $0x0,%edx
  3a:   48 85 c0                test   %rax,%rax
  3d:   48 0f 48 c2             cmovs  %rdx,%rax // note %edx is %rdx
```

A conditional move (`cmovs`) instruction is used here. That is, our optimized version is (on my machine) actually worse than the straightforward equivalent.

A second example is the function `clip_range()` above. It is compiled to

```
   0:   48 89 f8                mov    %rdi,%rax
   3:   48 29 f2                sub    %rsi,%rdx
   6:   31 c9                   xor    %ecx,%ecx
   8:   48 29 f0                sub    %rsi,%rax
   b:   78 0a                   js     17 <_Z2CLlll+0x17>  // the branch
   d:   48 39 d0                cmp    %rdx,%rax
  10:   48 89 d1                mov    %rdx,%rcx
  13:   48 0f 4e c8             cmovle %rax,%rcx
  17:   48 8d 04 0e             lea    (%rsi,%rcx,1),%rax
```

Now we replace the code by

```
1   inline long clip_range(long x, long mi, long ma)
2   {
3       x -= mi;
4       if ( x<0 )  x = 0;
5   //    else  // commented out to make (compiled) function really branchless
6       {
7           ma -= mi;
8           if ( x>ma )  x = ma;
9       }
10      x += mi;
11  }
```

Then the compiler generates branchless code:

```
   0:   48 89 f8                mov    %rdi,%rax
   3:   b9 00 00 00 00          mov    $0x0,%ecx
   8:   48 29 f0                sub    %rsi,%rax
   b:   48 0f 48 c1             cmovs  %rcx,%rax
   f:   48 29 f2                sub    %rsi,%rdx
  12:   48 39 d0                cmp    %rdx,%rax
  15:   48 0f 4f c2             cmovg  %rdx,%rax
  19:   48 01 f0                add    %rsi,%rax
```

Still, with CPUs that do not have a conditional move instruction (or some branchless equivalent of it) the techniques shown in this section can be useful.

1.12 Bit-wise rotation of a word

Neither C nor C++ have a statement for bit-wise rotation of a binary word (which may be considered a missing feature). The operation can be 'emulated' via [FXT: bits/bitrotate.h]:

```
1   static inline ulong bit_rotate_left(ulong x, ulong r)
2   // Return word rotated r bits to the left
3   // (i.e. toward the most significant bit)
4   {
5       return  (x<<r) | (x>>(BITS_PER_LONG-r));
6   }
```

As already mentioned, GCC emits exactly the CPU instruction that is *meant* here, even with non-constant argument `r`. Well done, GCC folks! Explicit use of the corresponding assembler instruction should not do any harm:

```
1   static inline ulong bit_rotate_right(ulong x, ulong r)
2   // Return word rotated r bits to the right
3   // (i.e. toward the least significant bit)
4   {
5   #if defined  BITS_USE_ASM    // use x86 asm code
6       return asm_ror(x, r);
7   #else
8       return  (x>>r) | (x<<(BITS_PER_LONG-r));
9   #endif
```

```
10  }
```

where we used [FXT: bits/bitasm-amd64.h]:

```
1   static inline ulong asm_ror(ulong x, ulong r)
2   {
3       asm ("rorq   %%cl, %0" : "=r" (x) : "0" (x), "c" (r));
4       return x;
5   }
```

Rotations using only a part of the word length are achieved by

```
1   static inline ulong bit_rotate_left(ulong x, ulong r, ulong ldn)
2   // Return ldn-bit word rotated r bits to the left
3   // (i.e. toward the most significant bit)
4   // Must have  0 <= r <= ldn
5   {
6       ulong m = ~0UL >> ( BITS_PER_LONG - ldn );
7       x &= m;
8       x = (x<<r) | (x>>(ldn-r));
9       x &= m;
10      return  x;
11  }
```

and

```
1   static inline ulong bit_rotate_right(ulong x, ulong r, ulong ldn)
2   // Return ldn-bit word rotated r bits to the right
3   // (i.e. toward the least significant bit)
4   // Must have  0 <= r <= ldn
5   {
6       ulong m = ~0UL >> ( BITS_PER_LONG - ldn );
7       x &= m;
8       x = (x>>r) | (x<<(ldn-r));
9       x &= m;
10      return  x;
11  }
```

Finally, the functions

```
1   static inline ulong bit_rotate_sgn(ulong x, long r, ulong ldn)
2   // Positive r --> shift away from element zero
3   {
4       if ( r > 0 )  return  bit_rotate_left(x, (ulong)r, ldn);
5       else          return  bit_rotate_right(x, (ulong)-r, ldn);
6   }
```

and (full-word version)

```
1   static inline ulong bit_rotate_sgn(ulong x, long r)
2   // Positive r --> shift away from element zero
3   {
4       if ( r > 0 )  return  bit_rotate_left(x, (ulong)r);
5       else          return  bit_rotate_right(x, (ulong)-r);
6   }
```

are sometimes convenient.

1.13 Functions related to bit-wise rotation and binary necklaces

We give several functions related to cyclic rotations of binary words and a class to generate binary necklaces.

1.13.1 Cyclic matching, minimum, and maximum

The following function determines whether there is a cyclic right shift of its second argument so that it matches the first argument. It is given in [FXT: bits/bitcyclic-match.h]:

```
1   static inline ulong bit_cyclic_match(ulong x, ulong y)
2   // Return  r if x==rotate_right(y, r) else return ~0UL.
3   // In other words: return
4   //   how often the right arg must be rotated right (to match the left)
5   // or, equivalently:
6   //   how often the left arg must be rotated left (to match the right)
7   {
8       ulong r = 0;
9       do
10      {
11          if ( x==y )  return r;
12          y = bit_rotate_right(y, 1);
13      }
14      while ( ++r < BITS_PER_LONG );
15
16      return ~0UL;
17  }
```

The functions shown work on the full length of the words, equivalents for the sub-word of the lowest `ldn` bits are given in the respective files. Just one example:

```
1   static inline ulong bit_cyclic_match(ulong x, ulong y, ulong ldn)
2   // Return  r if x==rotate_right(y, r, ldn) else return ~0UL
3   //  (using ldn-bit words)
4   {
5       ulong r = 0;
6       do
7       {
8           if ( x==y )  return r;
9           y = bit_rotate_right(y, 1, ldn);
10      }
11      while ( ++r < ldn );
12
13      return ~0UL;
14  }
```

The minimum among all cyclic shifts of a word can be computed via the following function given in [FXT: bits/bitcyclic-minmax.h]:

```
1   static inline ulong bit_cyclic_min(ulong x)
2   // Return minimum of all rotations of x
3   {
4       ulong r = 1;
5       ulong m = x;
6       do
7       {
8           x = bit_rotate_right(x, 1);
9           if ( x<m )  m = x;
10      }
11      while ( ++r < BITS_PER_LONG );
12
13      return  m;
14  }
```

1.13.2 Cyclic period and binary necklaces

Selecting from all n-bit words those that are equal to their cyclic minimum gives the sequence of the binary length-n necklaces, see chapter 17 on page 353. For example, with 6-bit words:

word	period	word	period
`.....1`	1	`..11.1`	6
`....i1`	6	`..1111`	6
`...ii1`	6	`.i.1.1`	2
`..iii1`	6	`.i.111`	6
`.iiii1`	6	`.ii.11`	3
`..i..1`	3	`.11111`	6
`..i.11`	6	`111111`	1

The values in each right column can be computed using [FXT: bits/bitcyclic-period.h]:

```
1   static inline ulong bit_cyclic_period(ulong x, ulong ldn)
2   // Return minimal positive bit-rotation that transforms x into itself.
3   //  (using ldn-bit words)
4   // The returned value is a divisor of ldn.
5   {
6       ulong y = bit_rotate_right(x, 1, ldn);
```

```
7       return  bit_cyclic_match(x, y, ldn) + 1;
8   }
```

It is possible to completely avoid the rotation of partial words: let d be a divisor of the word length n. Then the rightmost $(n-1)\,d$ bits of the word obtained as `x^(x>>d)` are zero exactly if the word has period d. Thereby we can use the following function body:

```
1       ulong sl = BITS_PER_LONG-ldn;
2       for (ulong s=1; s<ldn; ++s)
3       {
4           ++sl;
5           if ( 0==( (x^(x>>s)) << sl ) )  return s;
6       }
7       return ldn;
```

Testing for periods that are not divisors of the word length can be avoided as follows:

```
1       ulong f = tiny_factors_tab[ldn];
2       ulong sl = BITS_PER_LONG-ldn;
3       for (ulong s=1; s<ldn; ++s)
4       {
5           ++sl;
6           f >>= 1;
7           if ( 0==(f&1) )  continue;
8           if ( 0==( (x^(x>>s)) << sl ) )  return s;
9       }
10      return ldn;
```

The table of tiny factors used is shown in section 1.10 on page 24.

The version for `ldn==BITS_PER_LONG` can be optimized similarly:

```
1   static inline ulong bit_cyclic_period(ulong x)
2   // Return minimal positive bit-rotation that transforms x into itself.
3   // (same as bit_cyclic_period(x, BITS_PER_LONG) )
4   //
5   // The returned value is a divisor of the word length,
6   //   i.e. 1,2,4,8,...,BITS_PER_LONG.
7   {
8       ulong r = 1;
9       do
10      {
11          ulong y = bit_rotate_right(x, r);
12          if ( x==y )  return r;
13          r <<= 1;
14      }
15      while ( r < BITS_PER_LONG );
16
17      return  r;  // == BITS_PER_LONG
18  }
```

1.13.3 Generating all binary necklaces

We can generate all necklaces by the FKM algorithm given in section 17.1.1 on page 354. Here we specialize the method for binary words. The words generated are the cyclic maxima [FXT: `class bit_necklace` in bits/bit-necklace.h]:

```
1   class bit_necklace
2   {
3   public:
4       ulong a_;   // necklace
5       ulong j_;   // period of the necklace
6       ulong n2_;  // bit representing n: n2==2**(n-1)
7       ulong j2_;  // bit representing j: j2==2**(j-1)
8       ulong n_;   // number of bits in words
9       ulong mm_;  // mask of n ones
10      ulong tfb_; // for fast factor lookup
11
12  public:
13      bit_necklace(ulong n)  { init(n); }
14      ~bit_necklace()  { ; }
15
16      void init(ulong n)
```

```
17      {
18          if ( 0==n )  n = 1;  // avoid hang
19          if ( n>=BITS_PER_LONG )  n = BITS_PER_LONG;
20          n_ = n;
21
22          n2_ = 1UL<<(n-1);
23          mm_ = (~0UL) >> (BITS_PER_LONG-n);
24          tfb_ = tiny_factors_tab[n] >> 1;
25          tfb_ |= n2_;  // needed for n==BITS_PER_LONG
26          first();
27      }
28
29      void first()
30      {
31          a_ = 0;
32          j_ = 1;
33          j2_ = 1;
34      }
35
36      ulong data() const { return  a_; }
37      ulong period() const { return j_; }
```

The method for computing the successor is

```
1       ulong next()
2       // Create next necklace.
3       // Return the period, zero when current necklace is last.
4       {
5           if ( a_==mm_ )  { first();  return 0; }
6
7           do
8           {
9               // next lines compute index of highest zero, same result as
10              // j_ = highest_zero_idx( a_ ^ (~mm_)  );
11              // but the direct computation is faster:
12              j_ = n_ - 1;
13              ulong jb = 1UL << j_;
14              while ( 0!=(a_ & jb) )  { --j_;  jb>>=1; }
15
16              j2_ = 1UL << j_;
17              ++j_;
18              a_ |= j2_;
19              a_ = bit_copy_periodic(a_, j_, n_);
20          }
21          while ( 0==(tfb_ & j2_) );  // necklaces only
22
23          return  j_;
24      }
```

It uses the following function for periodic copying [FXT: bits/bitperiodic.h]:

```
1   static inline ulong bit_copy_periodic(ulong a, ulong p, ulong ldn)
2   // Return word that consists of the lowest p bits of a repeated
3   // in the lowest ldn bits (higher bits are zero).
4   // E.g.: if p==3, ldn=7 and a=*****xyz (8-bit), the return 0zxyzxyz.
5   // Must have p>0 and ldn>0.
6   {
7       a &= ( ~0UL >> (BITS_PER_LONG-p) );
8       for (ulong s=p; s<ldn; s<<=1)  { a |= (a<<s); }
9       a &= ( ~0UL >> (BITS_PER_LONG-ldn) );
10      return a;
11  }
```

Finally, we can easily detect whether a necklace is a Lyndon word:

```
1       ulong is_lyndon_word()  const  { return (j2_ & n2_); }
2
3       ulong next_lyn()
4       // Create next Lyndon word.
5       // Return the period (==n), zero when current necklace is last.
6       {
7           if ( a_==mm_ )  { first();  return 0; }
8           do  { next(); }  while ( !is_lyndon_word() );
9           return  n_;
10      }
11  };
```

About 54 million necklaces per second are generated (with $n = 32$), corresponding to a rate of 112 M/s for pre-necklaces [FXT: bits/bit-necklace-demo.cc].

1.13.4 Cyclic distance and cyclic XOR

A related function computes the cyclic distance between two words [FXT: bits/bitcyclic-dist.h]:

```
1   inline ulong bit_cyclic_dist(ulong a, ulong b)
2   // Return minimal bitcount of (t ^ b)
3   // where t runs through the cyclic rotations.
4   {
5       ulong d = ~0UL;
6       ulong t = a;
7       do
8       {
9           ulong z = t ^ b;
10          ulong e = bit_count( z );
11          if ( e < d )  d = e;
12          t = bit_rotate_right(t, 1);
13      }
14      while ( t!=a );
15      return  d;  // not reached
16  }
```

The functions [FXT: bits/bitcyclic-xor.h]

```
1   static inline ulong bit_cyclic_rxor(ulong x)
2   {
3       return x ^ bit_rotate_right(x, 1);
4   }
```

and

```
1   static inline ulong bit_cyclic_lxor(ulong x)
2   {
3       return x ^ bit_rotate_left(x, 1);
4   }
```

return a word where the number of bits is even. In order to produce a random value with an even number of set bits one can use either variant. If the bit count shall be odd, XOR the value with one[3] afterwards.

Iterated application always ends in a cycle, two examples using 6-bit words are:

```
.11111
i1....
1.1...  <--= cycle start
1111..
1...1.
11..11
.1.1.1
..1111  <--= cycle end
1.1...

.111..
.1..1.
.11.11  <--= cycle start
11.11.
1.11.1  <--= cycle end
.11.11
```

Zero is a fixed point (a period with cycle length one), a typical sequence using 8-bit words is:

```
.1..11.1
1...1.11
.1..111.
.11..1.1
11.111.1
..11..11
1.1.1.1.
11111111
........  <--= cycle start == cycle end
```

Cyclic shifts of a word produce cyclic shifts of the same cycle of words. A word and its complement produce the same result.

[3] Actually any value with an odd number of set bits will do for the XOR.

The inverse functions need no rotation at all, the inverse of `bit_cyclic_rxor()` is the inverse Gray code (see section 1.16 on page 40):

```
1   static inline ulong bit_cyclic_inv_rxor(ulong x)
2   // Return v so that bit_cyclic_rxor(v) == x.
3   {
4       return  inverse_gray_code(x);
5   }
```

The argument `x` must have an even number of bits. If this is the case then the lowest bit of the result is zero. The complement of the returned value is also an inverse of `bit_cyclic_rxor()`.

The inverse of `bit_cyclic_lxor()` is the inverse reversed code (see section 1.16.6 on page 45):

```
1   static inline ulong bit_cyclic_inv_lxor(ulong x)
2   // Return v so that bit_cyclic_lxor(v) == x.
3   {
4       return  inverse_rev_gray_code(x);
5   }
```

We do not need to mask out the lowest bit because for valid arguments (that have an even number of bits) the high bits of the result are zero. This function can be used to solve the quadratic equation $v^2 + v = x$ in the finite field GF(2^n) when normal bases are used, see page 891.

1.14 Reversing the bits of a word

The bits of a binary word can efficiently be reversed by a sequence of steps that reverse the order of certain blocks. For 16-bit words, we need $4 = \log_2(16)$ such steps [FXT: bits/revbin-steps-demo.cc]:

```
[ 0 1 2 3 4 5 6 7 8 9 a b c d e f ]
[ 1 0 3 2 5 4 7 6 9 8 b a d c f e ]  <--= pairs swapped
[ 3 2 1 0 7 6 5 4 b a 9 8 f e d c ]  <--= groups of 2 swapped
[ 7 6 5 4 3 2 1 0 f e d c b a 9 8 ]  <--= groups of 4 swapped
[ f e d c b a 9 8 7 6 5 4 3 2 1 0 ]  <--= groups of 8 swapped
```

1.14.1 Swapping adjacent bit blocks

We need a couple of auxiliary functions given in [FXT: bits/bitswap.h]. Pairs of adjacent bits can be swapped via

```
1   static inline ulong bit_swap_1(ulong x)
2   // Return x with neighbour bits swapped.
3   {
4   #if  BITS_PER_LONG == 32
5       ulong m = 0x55555555UL;
6   #else
7   #if  BITS_PER_LONG == 64
8       ulong m = 0x5555555555555555UL;
9   #endif
10  #endif
11      return  ((x & m) << 1) | ((x & (~m)) >> 1);
12  }
```

The 64-bit branch is omitted in the following examples. Adjacent groups of 2 bits are swapped by

```
1   static inline ulong bit_swap_2(ulong x)
2   // Return x with groups of 2 bits swapped.
3   {
4       ulong m = 0x33333333UL;
5       return  ((x & m) << 2) | ((x & (~m)) >> 2);
6   }
```

Equivalently,

```
1   static inline ulong bit_swap_4(ulong x)
2   // Return x with groups of 4 bits swapped.
3   {
4       ulong m = 0x0f0f0f0fUL;
```

```
5       return  ((x & m) << 4) | ((x & (~m)) >> 4);
6   }
```

and

```
1   static inline ulong bit_swap_8(ulong x)
2   // Return x with groups of 8 bits swapped.
3   {
4       ulong m = 0x00ff00ffUL;
5       return  ((x & m) << 8) | ((x & (~m)) >> 8);
6   }
```

When swapping half-words (here for 32-bit architectures)

```
1   static inline ulong bit_swap_16(ulong x)
2   // Return x with groups of 16 bits swapped.
3   {
4       ulong m = 0x0000ffffUL;
5       return  ((x & m) << 16) | ((x & (m<<16)) >> 16);
6   }
```

GCC is clever enough to recognize that the whole operation is equivalent to a (left or right) word rotation and indeed emits just a single rotate instruction. We could also use the bit-rotate function from section 1.12 on page 28, or

```
return (x << 16) | (x >> 16);
```

1.14.2 Bit-reversing binary words

The shown functions are taken from [FXT: bits/revbin.h]. The following is a 64-bit version of `revbin()`

```
1   static inline ulong revbin(ulong x)
2   // Return x with reversed bit order.
3   {
4       x = bit_swap_1(x);
5       x = bit_swap_2(x);
6       x = bit_swap_4(x);
7       x = bit_swap_8(x);
8       x = bit_swap_16(x);
9   #if  BITS_PER_LONG >= 64
10      x = bit_swap_32(x);
11  #endif
12      return x;
13  }
```

For 32-bit machines the `bit_swap_32()` line would have to be omitted.

The steps after `bit_swap_4()` correspond to a byte-reverse operation. This operation is just one assembler instruction for many CPUs (‘`bswap`’). The inline assembler with GCC for AMD64 CPUs is given in [FXT: bits/bitasm-amd64.h]:

```
1   static inline ulong asm_bswap(ulong x)
2   {
3       asm ("bswap %0" : "=r" (x) : "0" (x));
4       return x;
5   }
```

We use it for byte reversion when available:

```
1   static inline ulong bswap(ulong x)
2   // Return word with reversed byte order.
3   {
4   #ifdef BITS_USE_ASM
5       x = asm_bswap(x);
6   #else
7       x = bit_swap_8(x);
8       x = bit_swap_16(x);
9   #if  BITS_PER_LONG >= 64
10      x = bit_swap_32(x);
11  #endif
12  #endif // def BITS_USE_ASM
13      return x;
14  }
```

The function actually used for bit reversion is good for both 32 and 64 bit words:

```
1   static inline ulong revbin(ulong x)
2   {
3       x = bit_swap_1(x);
4       x = bit_swap_2(x);
5       x = bit_swap_4(x);
6       x = bswap(x);
7       return x;
8   }
```

One can generate the masks in the process as follows:

```
1   static inline ulong revbin(ulong x)
2   {
3       ulong s = BITS_PER_LONG >> 1;
4       ulong m = ~0UL >> s;
5       while ( s )
6       {
7           x = ( (x & m) << s ) ^ ( (x & (~m)) >> s );
8           s >>= 1;
9           m ^= (m<<s);
10      }
11      return  x;
12  }
```

Note that the above function will not always beat the obvious, bit-wise algorithm:

```
1   static inline ulong revbin(ulong x)
2   {
3       ulong r = 0,  ldn = BITS_PER_LONG;
4       while ( ldn-- != 0 )
5       {
6           r <<= 1;
7           r += (x&1);
8           x >>= 1;
9       }
10      return  r;
11  }
```

Therefore the function

```
1   static inline ulong revbin(ulong x, ulong ldn)
2   // Return word with the ldn least significant bits
3   //   (i.e. bit_0 ... bit_{ldn-1})  of x reversed,
4   //   the other bits are set to zero.
5   {
6       return  revbin(x) >> (BITS_PER_LONG-ldn);
7   }
```

should only be used when `ldn` is not too small, else replaced by the trivial algorithm.

One can also use table lookups methods so that, for example, eight bits are reversed at a time using a 256-byte table. We give the routine for full words:

```
1   unsigned char revbin_tab[256]; // reversed 8-bit words
2   ulong revbin_t(ulong x)
3   {
4       ulong r = 0;
5       for (ulong k=0; k<BYTES_PER_LONG; ++k)
6       {
7           r <<= 8;
8           r |= revbin_tab[ x & 255 ];
9           x >>= 8;
10      }
11      return r;
12  }
```

The routine can be optimized by unrolling to avoid all branches:

```
1   static inline ulong revbin_t(ulong x)
2   {
3       ulong r      = revbin_tab[ x & 255 ];  x >>= 8;
4       r <<= 8;  r |= revbin_tab[ x & 255 ];  x >>= 8;
5       r <<= 8;  r |= revbin_tab[ x & 255 ];  x >>= 8;
6   #if BYTES_PER_LONG > 4
7       r <<= 8;  r |= revbin_tab[ x & 255 ];  x >>= 8;
8       r <<= 8;  r |= revbin_tab[ x & 255 ];  x >>= 8;
```

```
9       r <<= 8;  r |= revbin_tab[ x & 255 ];  x >>= 8;
10      r <<= 8;  r |= revbin_tab[ x & 255 ];  x >>= 8;
11  #endif
12      r <<= 8;  r |= revbin_tab[ x ];
13      return r;
14  }
```

However, reversing the first 2^{30} binary words with this routine takes (on a 64-bit machine) longer than with the routine using the `bit_swap_NN()` calls, see [FXT: bits/revbin-tab-demo.cc].

Bit-hacker's life would be easier if there was a CPU instruction for reversing a binary word.

1.14.3 Generating the bit-reversed words in order

If the bit-reversed words have to be generated in the (reversed) counting order then there is a significantly cheaper way to do the update [FXT: bits/revbin-upd.h]:

```
1   static inline ulong revbin_upd(ulong r, ulong h)
2   // Let n=2**ldn and h=n/2.
3   // Then, with r == revbin(x, ldn) at entry, return revbin(x+1, ldn)
4   // Note: routine will hang if called with r the all-ones word
5   {
6       while ( !((r^=h)&h) )  h >>= 1;
7       return  r;
8   }
```

Now assume we want to generate the bit-reversed words of all $N = 2^n$ words smaller than 2^n. The total number of branches with the `while`-loop can be estimated by observing that for half of the updates just one bit changes, for a quarter two bits change, three bits change for one eighth of all updates, and so on. Thereby the loop executes less than $2\,N$ times:

$$N\left(\frac{1}{2}+\frac{2}{4}+\frac{3}{8}+\frac{4}{16}+\cdots+\frac{\log_2(N)}{N}\right) \;=\; N\sum_{j=1}^{\log_2(N)}\frac{j}{2^j} \;<\; 2\,N \qquad (1.14\text{-}1)$$

Observing that the updates that involve a single bit change occur at every second step we can avoid half of all branches.

For large vales of N the following method can be significantly faster if a fast routine is available for the computation of the least significant bit in a word. The underlying observation is that for a fixed word of size n there are just n different patterns of bit-changes with incrementing. We generate a lookup table of the bit-reversed patterns, `utab[]`, an array of `BITS_PER_LONG` elements:

```
1   inline void make_revbin_upd_tab(ulong ldn)
2   // Initialize lookup table used by revbin_tupd()
3   {
4       utab[0] = 1UL<<(ldn-1);
5       for (ulong k=1; k<ldn; ++k)  utab[k] = utab[k-1] | (utab[k-1]>>1);
6   }
```

The change patterns for $n = 5$ start as

```
pattern  reversed pattern
  ...1     1....
  ..11     11...
  ...1     1....
  .111     111..
  ...1     1....
  ..11     11...
  ...1     1....
  1111     1111.
  ...1     1....
  ..11     11...
```

The crucial observation is that the pattern with x set bits is used for the update of k to $k+1$ when the lowest zero of k is at position $x-1$:

```
                      used when the lowest
              reversed  zero of k is at index:
utab[0]=       1....      0
utab[1]=       11...      1
utab[2]=       111..      2
utab[3]=       1111.      3
utab[4]=       11111      4
```

The update routine can now be implemented as

```
1   inline ulong revbin_tupd(ulong r, ulong k)
2   // Let r==revbin(k, ldn) then
3   // return revbin(k+1, ldn).
4   // NOTE 1: need to call make_revbin_upd_tab(ldn) before usage
5   //         where ldn=log_2(n)
6   // NOTE 2: different argument structure than revbin_upd()
7   {
8       k = lowest_one_idx(~k);  // lowest zero idx
9       r ^= utab[k];
10      return r;
11  }
```

The revbin-update routines are used for the revbin permutation described in section 2.1.

	30 bits	16 bits	8 bits	
Update, bit-wise	1.00	1.00	1.00	`revbin_upd()`
Update, table	0.99	1.08	1.15	`revbin_tupd()`
Full, masks	0.74	0.81	0.86	`revbin()`
Full, 8-bit table	1.77	1.94	2.06	`revbin_t()`
Full32, 8-bit table	0.83	0.90	0.96	`revbin_t_le32()`
Full16, 8-bit table	—	0.54	0.58	`revbin_t_le16()`
Full, generated masks	2.97	3.25	3.45	[page 36]
Full, bit-wise	8.76	5.77	2.50	[page 36]

Figure 1.14-A: Relative performance of the revbin-update and (full) revbin routines. The timing of the bit-wise update routine is normalized to one. Values in each column should be compared, smaller values correspond to faster routines. A column labeled "N bits" gives the timing for reversing the N least significant bits of a word.

The relative performance of the different revbin routines is shown in figure 1.14-A. As a surprise, the full-word revbin function is consistently faster than both of the update routines. This is mainly due to the fact that the machine used (see appendix B on page 913) has a byte swap instruction. As the performance of table lookups is highly machine dependent your results can be very different.

1.14.4 Alternative techniques for in-order generation

The following loop, due to Brent Lehmann [priv.comm.], also generates the bit-reversed words in succession:

```
1       ulong n = 32;  // a power of two
2       ulong p = 0, s = 0, n2 = 2*n;
3       do
4       {
5           // here: s is the bit-reversed word
6           p += 2;
7           s ^= n - (n / (p&-p));
8       }
9       while ( p<n2 );
```

The revbin-increment is branchless but involves a division which usually is an expensive operation. With a fast bit-scan function the loop should be replaced by

```
1       do
2       {
3           p += 1;
4           s ^= n - (n >> (lowest_one_idx(p)+1));
5       }
6       while ( p<n );
```

A recursive algorithm for the generation of the bit-reversed words in order is given in [FXT: bits/revbin-rec-demo.cc]:

```
1   ulong N;
```

```
2   void revbin_rec(ulong f, ulong n)
3   {
4       // visit( f )
5       for (ulong m=N>>1; m>n; m>>=1)  revbin_rec(f+m, m);
6   }
```

One has to call `revbin_rec(0, 0)` to generate all `N`-bit bit-reversed words.

A technique to generate all revbin pairs in a pseudo random order is given in section 39.3 on page 861.

1.15 Bit-wise zip

The bit-wise zip (bit-zip) operation moves the lower half bits to even indices and higher half bits to odd indices. For example, with 8-bit words the permutation of bits is (see section 2.5 on page 101):

```
[ a b c d A B C D ]  |-->  [ a A b B c C d D ]
```

A straightforward implementation is

```
1   ulong bit_zip(ulong a, ulong b)
2   {
3       ulong x = 0;
4       ulong m = 1, s = 0;
5       for (ulong k=0; k<(BITS_PER_LONG/2); ++k)
6       {
7           x |= (a & m) << s;
8           ++s;
9           x |= (b & m) << s;
10          m <<= 1;
11      }
12      return  x;
13  }
```

Its inverse (bit-unzip) moves even indexed bits to the lower half-word and odd indexed bits to the higher half-word:

```
1   void bit_unzip(ulong x, ulong &a, ulong &b)
2   {
3       a = 0;  b = 0;
4       ulong m = 1, s = 0;
5       for (ulong k=0; k<(BITS_PER_LONG/2); ++k)
6       {
7           a |= (x & m) >> s;
8           ++s;
9           m <<= 1;
10          b |= (x & m) >> s;
11          m <<= 1;
12      }
13  }
```

The optimized versions (see [FXT: bits/bitzip.h]), using ideas similar to those in `revbin()` and `bit_count()`, are

```
1   static inline ulong bit_zip(ulong x)
2   {
3   #if  BITS_PER_LONG == 64
4       x = butterfly_16(x);
5   #endif
6       x = butterfly_8(x);
7       x = butterfly_4(x);
8       x = butterfly_2(x);
9       x = butterfly_1(x);
10      return  x;
11  }
```

and

```
1   static inline ulong bit_unzip(ulong x)
2   {
3       x = butterfly_1(x);
4       x = butterfly_2(x);
5       x = butterfly_4(x);
6       x = butterfly_8(x);
```

```
7   #if  BITS_PER_LONG == 64
8       x = butterfly_16(x);
9   #endif
10      return  x;
11  }
```

Both use the `butterfly_*()`-functions which are defined in [FXT: bits/bitbutterfly.h]:

```
1   static inline ulong butterfly_4(ulong x)
2   {
3   #if  BITS_PER_LONG == 64
4       const ulong ml = 0x0f000f000f000f00UL;
5   #else
6       const ulong ml = 0x0f000f00UL;
7   #endif
8       const ulong s = 4;
9       const ulong mr = ml >> s;
10      const ulong t = ((x & ml) >> s ) | ((x & mr) << s );
11      x = (x & ~(ml | mr)) | t;
12      return  x;
13  }
```

Laszlo Hars suggests [priv.comm.] the following routine (version for 32-bit words), which can be obtained by making the compile-time constants explicit:

```
1   inline uint32 bit_zip(uint32 x)
2   {
3       x = ((x & 0x0000ff00) << 8) | ((x >> 8) & 0x0000ff00) | (x & 0xff0000ff);
4       x = ((x & 0x00f000f0) << 4) | ((x >> 4) & 0x00f000f0) | (x & 0xf00ff00f);
5       x = ((x & 0x0c0c0c0c) << 2) | ((x >> 2) & 0x0c0c0c0c) | (x & 0xc3c3c3c3);
6       x = ((x & 0x22222222) << 1) | ((x >> 1) & 0x22222222) | (x & 0x99999999);
7       return x;
8   }
```

Functions that zip/unzip the bits of (the lower half of) two words are

```
1   #define  BPLH  (BITS_PER_LONG/2)
2
3   static inline ulong bit_zip2(ulong x, ulong y)
4   // Two-word version:
5   // only the lower half of x and y are merged
6   {
7       return  bit_zip( (y<<BPLH) + x );
8   }
```

and

```
1   static inline void bit_unzip2(ulong t, ulong &x, ulong &y)
2   // Two-word version:
3   // only the lower half of x and y are filled
4   {
5       t = bit_unzip(t);
6       y = t >> BPLH;
7       x = t ^ (y<<BPLH);
8   }
```

A bit-zip version for words whose higher half is zero is given in [FXT: bits/bitzip0.h]:

```
1   static inline ulong bit_zip0(ulong x)
2   // Return word with lower half bits in even indices.
3   {
4       { ulong m = 0x00000000ffff0000UL;  m&=x;  x^=m;  x^=(m<<16); }
5       { ulong m = 0x0000ff000000ff00UL;  m&=x;  x^=m;  x^=(m<<8); }
6       { ulong m = 0x00f000f000f000f0UL;  m&=x;  x^=m;  x^=(m<<4); }
7       { ulong m = 0x0c0c0c0c0c0c0c0cUL;  m&=x;  x^=m;  x^=(m<<2); }
8       { ulong m = 0x2222222222222222UL;  m&=x;  x^=m;  x^=(m<<1); }
9       return  x;
10  }
```

1.16 Gray code and parity

The *Gray code* of a binary word can easily be computed by [FXT: bits/graycode.h]

```
1   static inline ulong gray_code(ulong x)
2   {
3       return  x ^ (x>>1);
4   }
```

```
       k:       bin(k)       g(k)     g(g(k))     g(2*k)   g(2*k+1)
       0:       ......    ......    ......    ......    .....1
       1:       .....1    .....1    .....1    ....11    ....1.
       2:       ....1.    ....11    ....1.    ...11.    ...111
       3:       ....11    ....1.    ....11    ...1.1    ...1..
       4:       ...1..    ...11.    ...1.1    ..11..    ..11.1
       5:       ...1.1    ...111    ...1..    ..1111    ..111.
       6:       ...11.    ...1.1    ...111    ..1.1.    ..1.11
       7:       ...111    ...1..    ...11.    ..1..1    ..1...
       8:       ..1...    ..11..    ..1.1.    .11...    .11..1
       9:       ..1..1    ..11.1    ..1.11    .11.11    .11.1.
      10:       ..1.1.    ..1111    ..1...    .1111.    .11111
      11:       ..1.11    ..111.    ..1..1    .111.1    .111..
      12:       ..11..    ..1.1.    ..1111    .1.1..    .1.1.1
      13:       ..11.1    ..1.11    ..111.    .1.111    .1.11.
      14:       ..111.    ..1..1    ..11.1    .1..1.    .1..11
      15:       ..1111    ..1...    ..11..    .1...1    .1....
      16:       .1....    .11...    .1.1..    11....    11...1
      17:       .1...1    .11..1    .1.1.1    11..11    11..1.
      18:       .1..1.    .11.11    .1.11.    11.11.    11.111
      19:       .1..11    .11.1.    .1.111    11.1.1    11.1..
      20:       .1.1..    .1111.    .1...1    1111..    1111.1
      21:       .1.1.1    .11111    .1....    111111    11111.
      22:       .1.11.    .111.1    .1..11    111.1.    111.11
      23:       .1.111    .111..    .1..1.    111..1    111...
      24:       .11...    .1.1..    .1111.    1.1...    1.1..1
      25:       .11..1    .1.1.1    .11111    1.1.11    1.1.1.
      26:       .11.1.    .1.111    .111..    1.111.    1.1111
      27:       .11.11    .1.11.    .111.1    1.11.1    1.11..
      28:       .111..    .1..1.    .11.11    1..1..    1..1.1
      29:       .111.1    .1..11    .11.1.    1..111    1..11.
      30:       .1111.    .1...1    .11..1    1...1.    1...11
      31:       .11111    .1....    .11...    1....1    1.....
```

Figure 1.16-A: Binary words, their Gray code, squared Gray code, and Gray codes of even and odd values.

Gray codes of consecutive values differ in one bit. Squared Gray codes of consecutive values differ in one or two bits. Gray codes of values that have a difference of a power of two differ in two bits. Gray codes of even/odd values have an even/odd number of bits set, respectively. This is demonstrated in [FXT: bits/gray2-demo.cc], whose output is given in figure 1.16-A.

In order to produce a random value with an even/odd number of bits set, set the lowest bit of a random number to zero/one, respectively, and take the Gray code.

Computing the inverse Gray code is slightly more expensive. Understanding the Gray code as 'bit-wise difference modulo 2' leads to the idea of computing the 'bit-wise sums modulo 2' for the inverse:

```
1   static inline ulong inverse_gray_code(ulong x)
2   {
3       // VERSION 1 (integration modulo 2):
4       ulong h=1, r=0;
5       do
6       {
7           if ( x & 1 )  r^=h;
8           x >>= 1;
9           h = (h<<1)+1;
10      }
11      while ( x!=0 );
12      return r;
13  }
```

For n-bit words, n-fold application of the Gray code gives back the original word. Using the symbol G for the Gray code (operator) we have $G^n = \mathrm{id}$, so $G^{n-1} \circ G = \mathrm{id} = G^{-1} \circ G$. That is, applying the Gray code computation $n-1$ times gives the inverse Gray code. Thus we can simplify to

```
1       // VERSION 2 (apply graycode BITS_PER_LONG-1 times):
2       ulong r = BITS_PER_LONG;
3       while ( --r )  x ^= x>>1;
4       return x;
```

Applying the Gray code twice is identical to `x^=x>>2;`, applying it four times is `x^=x>>4;`, and the idea holds for all powers of two. This leads to the most efficient way to compute the inverse Gray code:

```
1       // VERSION 3 (use: gray ** BITSPERLONG == id):
2       x ^= x>>1;  // gray ** 1
3       x ^= x>>2;  // gray ** 2
4       x ^= x>>4;  // gray ** 4
5       x ^= x>>8;  // gray ** 8
6       x ^= x>>16; // gray ** 16
7       // here: x = gray**31(input)
8       // note: the statements can be reordered at will
9   #if  BITS_PER_LONG >= 64
10      x ^= x>>32;  // for 64bit words
11  #endif
12      return  x;
```

1.16.1 The parity of a binary word

The *parity* of a word is its bit-count modulo two. The inverse Gray code of a word contains at each bit position the parity of all bits of the input left from it (including itself). Thereby we use the lowest bit [FXT: bits/parity.h]:

```
1   static inline ulong parity(ulong x)
2   // return 1 if the number of set bits is even, else 0
3   {
4       return  inverse_gray_code(x) & 1;
5   }
```

Be warned that the parity bit of many CPUs is the complement of the above. With the x86-architecture the parity bit only takes in account the lowest byte, therefore [FXT: bits/bitasm-i386.h]:

```
1   static inline ulong asm_parity(ulong x)
2   {
3       x ^= (x>>16);
4       x ^= (x>>8);
5       asm ("addl  $0, %0  \n"
6            "setnp %%al    \n"
7            "movzx %%al, %0"
8            : "=r" (x) : "0" (x) : "eax");
9       return x;
10  }
```

The equivalent code for the AMD64 CPU is [FXT: bits/bitasm-amd64.h]:

```
1   static inline ulong asm_parity(ulong x)
2   {
3       x ^= (x>>32);
4       x ^= (x>>16);
5       x ^= (x>>8);
6       asm ("addq  $0, %0  \n"
7            "setnp %%al    \n"
8            "movzx %%al, %0"
9            : "=r" (x) : "0" (x) : "eax");
10      return x;
11  }
```

1.16.2 Byte-wise Gray code and parity

A byte-wise Gray code can be computed using (32-bit version)

```
1   static inline ulong byte_gray_code(ulong x)
2   // Return the Gray code of bytes in parallel
3   {
4       return  x ^ ((x & 0xfefefefe)>>1);
5   }
```

Its inverse is

```
1   static inline ulong byte_inverse_gray_code(ulong x)
2   // Return the inverse Gray code of bytes in parallel
```

```
3   {
4       x ^= ((x & 0xfefefefeUL)>>1);
5       x ^= ((x & 0xfcfcfcfcUL)>>2);
6       x ^= ((x & 0xf0f0f0f0UL)>>4);
7       return  x;
8   }
```

Thereby

```
1   static inline ulong byte_parity(ulong x)
2   // Return the parities of bytes in parallel
3   {
4       return  byte_inverse_gray_code(x) & 0x01010101UL;
5   }
```

1.16.3 Incrementing (counting) in Gray code

```
        k:     g(k)      g(2*k)      g(k) p     diff p    set
        0:   ......    ......    ...... .   ...... .    {}
        1:   ......1   .....11   .....1 1   .....+ 1    {0}
        2:   .....11   ....11.   ....11 .   ....+1 .    {0, 1}
        3:   .....1.   ....1.1   ....1. 1   ....1- 1    {1}
        4:   ....11.   ...11..   ...11. .   ...+1. .    {1, 2}
        5:   ....111   ...1111   ...111 1   ...11+ 1    {0, 1, 2}
        6:   ....1.1   ...1.1.   ...1.1 .   ...1-1 .    {0, 2}
        7:   ....1..   ...1..1   ...1.. 1   ...1.- 1    {2}
        8:   ...11..   ..11...   ..11.. .   ..+1.. .    {2, 3}
        9:   ...11.1   ..11.11   ..11.1 1   ..11.+ 1    {0, 2, 3}
       10:   ...1111   ..1111.   ..1111 .   ..11+1 .    {0, 1, 2, 3}
       11:   ...111.   ..111.1   ..111. 1   ..111- 1    {1, 2, 3}
       12:   ...1.1.   ..1.1..   ..1.1. .   ..1-1. .    {1, 3}
       13:   ...1.11   ..1.111   ..1.11 1   ..1.1+ 1    {0, 1, 3}
       14:   ...1..1   ..1..1.   ..1..1 .   ..1.-1 .    {0, 3}
       15:   ...1...   ..1...1   ..1... 1   ..1..- 1    {3}
       16:   ..11...   .11....   .11... .   .+1... .    {3, 4}
       17:   ..11..1   .11..11   .11..1 1   .11..+ 1    {0, 3, 4}
```

Figure 1.16-B: The Gray code equals the Gray code of doubled value shifted to the right once. Equivalently, we can separate the lowest bit which equals the parity of the other bits. The last column shows that the changes with each increment always happen one position left of the rightmost bit.

Let $g(k)$ be the Gray code of a number k. We are interested in efficiently generating $g(k+1)$. Using the observation shown in figure 1.16-B we can implement a fast Gray counter if we use a spare bit to keep track of the parity of the Gray code word. The following routine does this [FXT: bits/nextgray.h]:

```
1   inline ulong next_gray2(ulong x)
2   // With input x==gray_code(2*k) the return is gray_code(2*k+2).
3   // Let x1 be the word x shifted right once
4   // and i1 its inverse Gray code.
5   // Let r1 be the return r shifted right once.
6   // Then r1 = gray_code(i1+1).
7   // That is, we have a Gray code counter.
8   // The argument must have an even number of bits.
9   {
10      x ^= 1;
11      x ^= (lowest_one(x) << 1);
12      return x;
13  }
```

To obtain a Gray counter, start with `x=0`, increment with `x=next_gray2(pg)` and use the words `g=x>>1`:

```
1       ulong x = 0;
2       for (ulong k=0; k<n2; ++k)
3       {
4           ulong g = x>>1;
5           x = next_gray2(x);
6           // here:  g == gray_code(k);
7       }
```

8

This is shown in [FXT: bits/bit-nextgray-demo.cc].

To start at an arbitrary (Gray code) value `g` compute

```
x = (g<<1) ^ parity(g)
```

in order to use the statement `x=next_gray2(x)` for later increments.

If one works with a set whose elements are the set bits in the Gray code then the parity is the set size k modulo two. The increment can then be achieved as follows: if k is even then, if the first element is zero, then remove it, else prepend the element zero. If k is odd then, if the first element equals the second minus one, then remove the second element, else insert at the second position the element equal to the first element plus one. Further, the decrement is obtained by simply swapping the actions for even and odd parity.

If one works with an array that contains the elements of the set it is more convenient to actually do the described operations at the end of the array. This leads to the (loopless) algorithm for subsets in minimal-change order that is given in section 8.2 on page 203.

1.16.4 The Thue-Morse sequence

The sequence of parities of the binary words,

```
0110100110010110100101100110100110010110011010011...
```

is called the *Thue-Morse sequence* (entry A010060 of [240]). It appears in various seemingly unrelated contexts, see [8] and section 36.1 on page 713.

The sequence can be generated with [FXT: `class thue_morse` in bits/thue-morse.h]:

```
class thue_morse
// Thue-Morse sequence
{
public:
    ulong k_;
    ulong tm_;

public:
    thue_morse(ulong k)  { init(k); }
    ~thue_morse()  { ; }

    ulong init(ulong k)
    {
        k_ = k;
        tm_ = parity(k_);
        return tm_;
    }

    ulong data()  { return tm_; }

    ulong next()
    {
        ulong x = k_ ^ (k_ + 1);
        ++k_;
        x ^= x>>1;        // highest bit that changed with increment
        x &= 0x5555555555555555UL;  // 64-bit version
        tm_ ^= ( x!=0 );  // change if highest changed bit was at even index
        return tm_;
    }
};
```

The rate of generation is about 435 million per second (5 cycles per update) [FXT: bits/thue-morse-demo.cc].

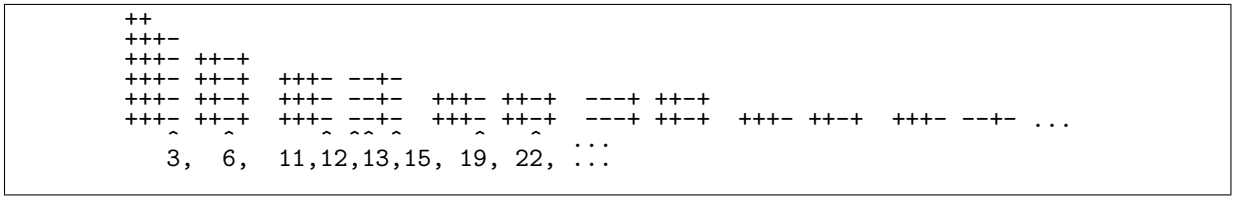

Figure 1.16-C: A construction for the GRS sequence.

1.16.5 The Golay-Rudin-Shapiro sequence *

The function [FXT: bits/grsnegative.h]

```
1   static inline ulong grs_negative_q(ulong x)
2   {
3       return  parity( x & (x>>1) );
4   }
```

returns one for indices where the *Golay-Rudin-Shapiro sequence* (or *GRS sequence*) has a negative value. The function returns one for `x` in the sequence

```
3, 6, 11, 12, 13, 15, 19, 22, 24, 25, 26, 30, 35, 38, 43, 44, 45,
47, 48, 49, 50, 52, 53, 55, 59, 60, 61, 63, 67, 70, 75, 76, 77,
79, 83, 86, 88, 89, 90, 94, 96, 97, 98, 100, 101, 103, 104, 105,
106, 110, 115, 118, 120, 121, 122, 126, 131, 134, 139, 140, ...
```

This is sequence A020985 of [240], see also section 36.3 on page 718.

The sequence can be obtained by starting with a sequence of two ones and in each step appending the left half and the negated right half of the values so far, see figure 1.16-C.

The algorithm counts the bit-pairs modulo 2. Note that the sequence `[1111]` contains three bit-pairs: `[11..]`, `[.11.]` and `[..11]`. The function proves to be useful in specialized versions of the fast Fourier- and Walsh transform, see section 22.4 on page 453.

1.16.6 The reversed Gray code

```
------------------------------------------------------------
111.1111...1111.............  = 0xef0f0000  == word
1..11...1...1...1.............. = gray_code
..11...1...1...1.............. = rev_gray_code
1.11.1.11111.1.1111111111111111 = inverse_gray_code
1.1..1.1.....1.1.............. = inverse_rev_gray_code
------------------------------------------------------------
...1....1111....111111111111111  = 0x10f0ffff  == word
...11...1...1...1.............. = gray_code
..11...1...1...1..............1 = rev_gray_code
...11111.1.11111.1.1.1.1.1.1.1.1 = inverse_gray_code
1111.....1.1.....1.1.1.1.1.1.1.1 = inverse_rev_gray_code
------------------------------------------------------------
......1.........................  = 0x2000000  == word
......11........................ = gray_code
.....11......................... = rev_gray_code
......11111111111111111111111111 = inverse_gray_code
1111111......................... = inverse_rev_gray_code
------------------------------------------------------------
111111.11111111111111111111111  = 0xfdffffff  == word
1.....11........................ = gray_code
.....11.......................1 = rev_gray_code
1.1.1..1.1.1.1.1.1.1.1.1.1.1.1 = inverse_gray_code
1.1.1.11.1.1.1.1.1.1.1.1.1.1.1 = inverse_rev_gray_code
------------------------------------------------------------
```

Figure 1.16-D: Four examples of the Gray code, reversed Gray code and their inverses with 32-bit words.

We define the *reversed Gray code* to be the bit-reversed word of the Gray code of the bit-reversed word. That is,

```
rev_gray_code(x) := revbin(gray_code(revbin(x)))
```

It turns out that the corresponding functions are identical to the Gray code versions up to the reversed shift operations (C-language operators '>>' replaced by '<<'). Thereby, computing the reversed Gray code is as easy as [FXT: bits/revgraycode.h]:

```
1  static inline ulong rev_gray_code(ulong x)
2  {
3      return  x ^ (x<<1);
4  }
```

Its inverse is

```
1  static inline ulong inverse_rev_gray_code(ulong x)
2  {
3      // use: rev_gray ** BITSPERLONG == id:
4      x ^= x<<1;  // rev_gray ** 1
5      x ^= x<<2;  // rev_gray ** 2
6      x ^= x<<4;  // rev_gray ** 4
7      x ^= x<<8;  // rev_gray ** 8
8      x ^= x<<16; // rev_gray ** 16
9      // here: x = rev_gray**31(input)
10     // note: the statements can be reordered at will
11 #if  BITS_PER_LONG >= 64
12     x ^= x<<32;  // for 64bit words
13 #endif
14     return  x;
15 }
```

Some examples with 32-bit words are shown in figure 1.16-D. The inverse reversed Gray code contains at each bit position the parity of all bits of the input right from it, including the bit itself. Especially, the word parity can be found in the highest bit of the inverse reversed Gray code.

The reversed Gray code preserves the lowest set bit while the Gray code preserves the highest.

Let G^{-1} and E^{-1} be the inverse Gray- and reversed Gray code of X, respectively. Then the bit-wise sum (XOR) of G^{-1} and E^{-1} equals X if the parity of X is zero, else it equals the complement X.

We note that taking the reversed Gray code of a binary word corresponds to multiplication with the binary polynomial $x+1$, and the inverse reversed Gray code is a method for fast exact division by $x+1$, see section 38.1.6 on page 820.

1.17 Bit sequency

Functions concerned with the sequency (number of zero-one transitions) are given in [FXT: bits/bitsequency.h]. Sequency counting:

```
1  static inline ulong bit_sequency(ulong x)
2  {
3      return bit_count( gray_code(x) );
4  }
```

The function assumes that all bits to the left of the word are zero, and all bits to right are equal to the lowest bit. For example, the sequency of the 8-bit word `[00011111]` is one. To take the lowest bit into account, add it to the sequency (then all sequencies are even).

Computation of the minimal binary word with given sequency:

```
1  static inline ulong first_sequency(ulong k)
2  // Return the first (i.e. smallest) word with sequency k,
3  // e.g.  00..00010101010 (seq 8)
4  // e.g.  00..00101010101 (seq 9)
5  // Must have:  0 <= k <= BITS_PER_LONG
6  {
7      return inverse_gray_code( first_comb(k) );
```

```
8   }
```

A faster version is (32-bit branch only):

```
1       if ( k==0 )  return 0;
2       const ulong m = 0xaaaaaaaaUL;
3       return  m >> (BITS_PER_LONG-k);
```

Computation of the maximal binary word with given sequency:

```
1   static inline ulong last_sequency(ulong k)
2   // Return the last (i.e. biggest) word with sequency k.
3   {
4       return inverse_gray_code( last_comb(k) );
5   }
```

The functions `first_comb(k)` and `last_comb(k)` return a word with k bits set at the low and high end, respectively (see section 1.26 on page 66).

Generation of all words with a given sequency, starting with the smallest, can be achieved with a function that computes the next word with the same sequency:

```
1   static inline ulong next_sequency(ulong x)
2   // Return smallest integer with highest bit at greater or equal
3   // position than the highest bit of x that has the same number
4   // of zero-one transitions (sequency) as x.
5   // The value of the lowest bit is conserved.
6   //
7   // Zero is returned when there is no further sequence.
8   {
9       x = gray_code(x);
10      x = next_colex_comb(x);
11      x = inverse_gray_code(x);
12      return x;
13  }
```

The inverse function, returning the previous word with the same sequency, is:

```
1   static inline ulong prev_sequency(ulong x)
2   {
3       x = gray_code(x);
4       x = prev_colex_comb(x);
5       x = inverse_gray_code(x);
6       return x;
7   }
```

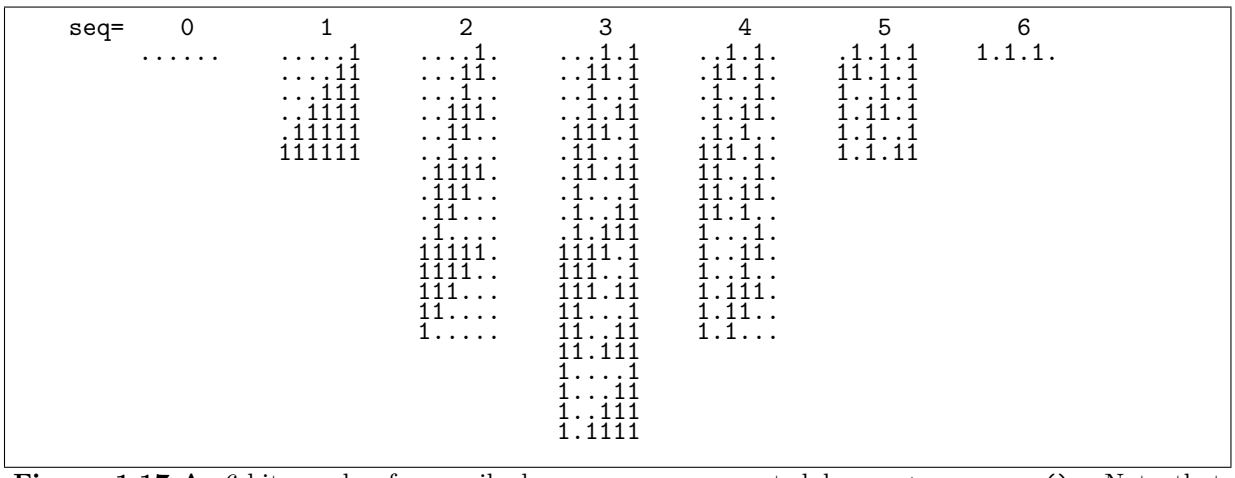

Figure 1.17-A: 6-bit words of prescribed sequency as generated by `next_sequency()`. Note that the transition at the lower end is not counted. This is consistent with sequency counting function `bit_sequency()`.

The list of all 6-bit words ordered by sequency is shown in figure 1.17-A. It was created with the program [FXT: bits/bitsequency-demo.cc].

We note that the sequency of a word can be 'complemented' as follows (32-bit version):

```
1   static inline ulong complement_sequency(ulong x)
2   // Return word whose sequency is BITS_PER_LONG - s
3   // where s is the sequency of x
4   {
5       return x ^ 0xaaaaaaaaUL;
6   }
```

1.18 Powers of the Gray code

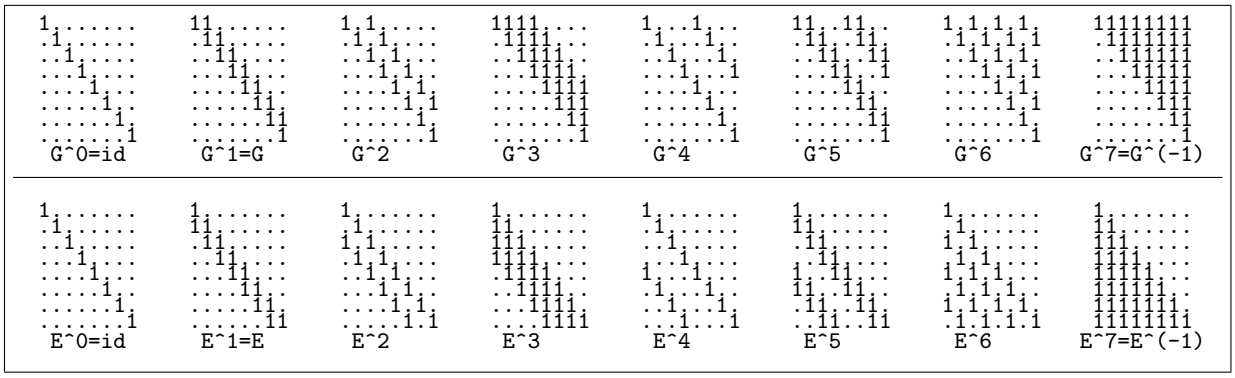

Figure 1.18-A: Powers of the matrices for the Gray code (top) and the reversed Gray code (bottom).

The Gray code is a bit-wise linear transform of a binary word. The 2^k-th power of the Gray code of x can be computed as `x ^ (x>>k)`. The e-th power can be computed as the bit-wise sum of the powers corresponding to the bits in the exponent. This motivates [FXT: bits/graypower.h]:

```
1   inline ulong gray_pow(ulong x, ulong e)
2   // Return (gray_code**e)(x)
3   // gray_pow(x, 1) == gray_code(x)
4   // gray_pow(x, BITS_PER_LONG-1) == inverse_gray_code(x)
5   {
6       e &= (BITS_PER_LONG-1);  // modulo BITS_PER_LONG
7       ulong s = 1;
8       while ( e )
9       {
10          if ( e & 1 )  x ^= x >> s;  // gray ** s
11          s <<= 1;
12          e >>= 1;
13      }
14      return  x;
15  }
```

The Gray code $g = [g_0, g_1, \ldots, g_7]$ of a 8-bit binary word $x = [x_0, x_1, \ldots, x_7]$ can be expressed as a matrix multiplication over GF(2) (dots for zeros):

```
  g    =        G          x
[g0]       [ 11...... ]  [x0]
[g1]       [ .11..... ]  [x1]
[g2]       [ ..11.... ]  [x2]
[g3]  =    [ ...11... ]  [x3]
[g4]       [ ....11.. ]  [x4]
[g5]       [ .....11. ]  [x5]
[g6]       [ ......11 ]  [x6]
[g7]       [ .......1 ]  [x7]
```

The powers of the Gray code correspond to multiplication with powers of the matrix G, shown in figure 1.18-A (bottom). The powers of the inverse Gray code for N-bit words (where N is a power of two) can be computed by the relation $G^e\, G^{N-e} = G^N = \mathrm{id}$.

```
1   inline ulong inverse_gray_pow(ulong x, ulong e)
2   // Return (inverse_gray_code**(e))(x)
```

```
3   //   == (gray_code**(-e))(x)
4   // inverse_gray_pow(x, 1) == inverse_gray_code(x)
5   // inverse_gray_pow(x, BITS_PER_LONG-1) == gray_code(x)
6   {
7       return  gray_pow(x, -e);
8   }
```

The matrices corresponding to the powers of the reversed Gray code are shown in figure 1.18-A (bottom). We just have to reverse the shift operator in the functions:

```
1   inline ulong rev_gray_pow(ulong x, ulong e)
2   // Return (rev_gray_code**e)(x)
3   {
4       e &= (BITS_PER_LONG-1);  // modulo BITS_PER_LONG
5       ulong s = 1;
6       while ( e )
7       {
8           if ( e & 1 )  x ^= x << s;  // rev_gray ** s
9           s <<= 1;
10          e >>= 1;
11      }
12      return  x;
13  }
```

The inverse function is

```
1   inline ulong inverse_rev_gray_pow(ulong x, ulong e)
2   // Return (inverse_rev_gray_code**(e))(x)
3   {
4       return  rev_gray_pow(x, -e);
5   }
```

1.19 Invertible transforms on words

The functions presented in this section are invertible 'transforms' on binary words. The names are chosen as 'some code', emphasizing the result of the transforms, similar to the convention used with the name 'Gray code'. The functions are given in [FXT: bits/bittransforms.h].

In the transform (*blue code*)

```
1   inline ulong blue_code(ulong a)
2   {
3       ulong s = BITS_PER_LONG >> 1;
4       ulong m = ~0UL << s;
5       do
6       {
7           a ^= ( (a&m) >> s );
8           s >>= 1;
9           m ^= (m>>s);
10      }
11      while ( s );
12      return  a;
13  }
```

the masks 'm' are (32-bit binary)

```
1111111111111111................
11111111........11111111........
1111....1111....1111....1111....
11..11..11..11..11..11..11..11..
1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.
```

The same masks are used in the *yellow code*

```
1   inline ulong yellow_code(ulong a)
2   {
3       ulong s = BITS_PER_LONG >> 1;
4       ulong m = ~0UL >> s;
5       do
6       {
7           a ^= ( (a&m) << s );
8           s >>= 1;
9           m ^= (m<<s);
10      }
```

```
11      while ( s );
12      return  a;
13  }
```

Both involve a computational work $\sim \log_2(b)$ where b is the number of bits per word (`BITS_PER_LONG`). The `blue_code` can be used as a fast implementation for the composition of a binary polynomial with $x+1$, see page 838. Note the names 'blue code' etc. are ad hoc terminology and not standard.

```
            blue                yellow
    0:     ......  0*      ................................  0
    1:     .....1  1*      1111111111111111111111111111111. 32
    2:     ....11  2       1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1. 16
    3:     ....1.  1       .1.1.1.1.1.1.1.1.1.1.1.1.1.1.1. 16
    4:     ...1.1  2       11..11..11..11..11..11..11..11.. 16
    5:     ...1..  1       ..11..11..11..11..11..11..11..11 16
    6:     ...11.  2*      .11..11..11..11..11..11..11..11. 16
    7:     ...111  3*      1..11..11..11..11..11..11..11..1 16
    8:     ..1111  4       1...1...1...1...1...1...1...1... 8
    9:     ..111.  3       .111.111.111.111.111.111.111.111 24
   10:     ..11..  2       ..1...1...1...1...1...1...1...1. 8
   11:     ..11.1  3       11.111.111.111.111.111.111.111.1 24
   12:     ..1.1.  2       .1...1...1...1...1...1...1...1.. 8
   13:     ..1.11  3       1.111.111.111.111.111.111.111.11 24
   14:     ..1..1  2       111.111.111.111.111.111.111.111. 24
   15:     ..1...  1       ...1...1...1...1...1...1...1...1 8
   16:     .1...1  2       1111....1111....1111....1111.... 16
   17:     .1....  1       ....1111....1111....1111....1111 16
   18:     .1..1.  2*      .1.11.1..1.11.1..1.11.1..1.11.1. 16
   19:     .1..11  3*      1.1..1.11.1..1.11.1..1.11.1..1.1 16
   20:     .1.1..  2*      ..1111....1111....1111....1111.. 16
   21:     .1.1.1  3*      11....1111....1111....1111....11 16
   22:     .1.111  4       1..1.11.1..1.11.1..1.11.1..1.11. 16
   23:     .1.11.  3       .11.1..1.11.1..1.11.1..1.11.1..1 16
   24:     .1111.  4       .1111....1111....1111....1111... 16
   25:     .11111  5       1....1111....1111....1111....111 16
   26:     .111.1  4       11.1..1.11.1..1.11.1..1.11.1..1. 16
   27:     .111..  3       ..1.11.1..1.11.1..1.11.1..1.11.1 16
   28:     .11.11  4       1.11.1..1.11.1..1.11.1..1.11.1.. 16
   29:     .11.1.  3       .1..1.11.1..1.11.1..1.11.1..1.11 16
   30:     .11...  2       ...1111....1111....1111....1111. 16
   31:     .11..1  3       111....1111....1111....1111....1 16
```

Figure 1.19-A: Blue and yellow transforms of the binary words 0,1,...,32. Bit-counts are shown at the right of each column. Fixed points are marked with asterisks.

The transforms of the binary words up to 31 are shown in figure 1.19-A, the lists were created with the program [FXT: bits/bittransforms-blue-demo.cc]. The parity of $B(a)$ is equal to the lowest bit of a. Up to the $a = 47$ the bit-count varies by ± 1 between successive values of $B(a)$, the transition $B(47) \to B(48)$ changes the bit-count by 3. The sequence of the indices a where the bit-count changes by more than one is

```
47, 51, 59, 67, 75, 79, 175, 179, 187, 195, 203, 207, 291, 299, 339, 347, 419, 427, ...
```

The yellow code might be a good candidate for 'randomization' of binary words. The blue code maps any range $[0 \ldots 2^k - 1]$ onto itself. Both the blue code and the yellow code are involutions (self-inverse).

The transforms (*red code*)

```
1   inline ulong red_code(ulong a)
2   {
3       ulong s = BITS_PER_LONG >> 1;
4       ulong m = ~0UL >> s;
5       do
6       {
7           ulong u = a & m;
8           ulong v = a ^ u;
9           a = v ^ (u<<s);
10          a ^= (v>>s);
11          s >>= 1;
12          m ^= (m<<s);
13      }
14      while ( s );
15      return  a;
16  }
```

and (*green code*)

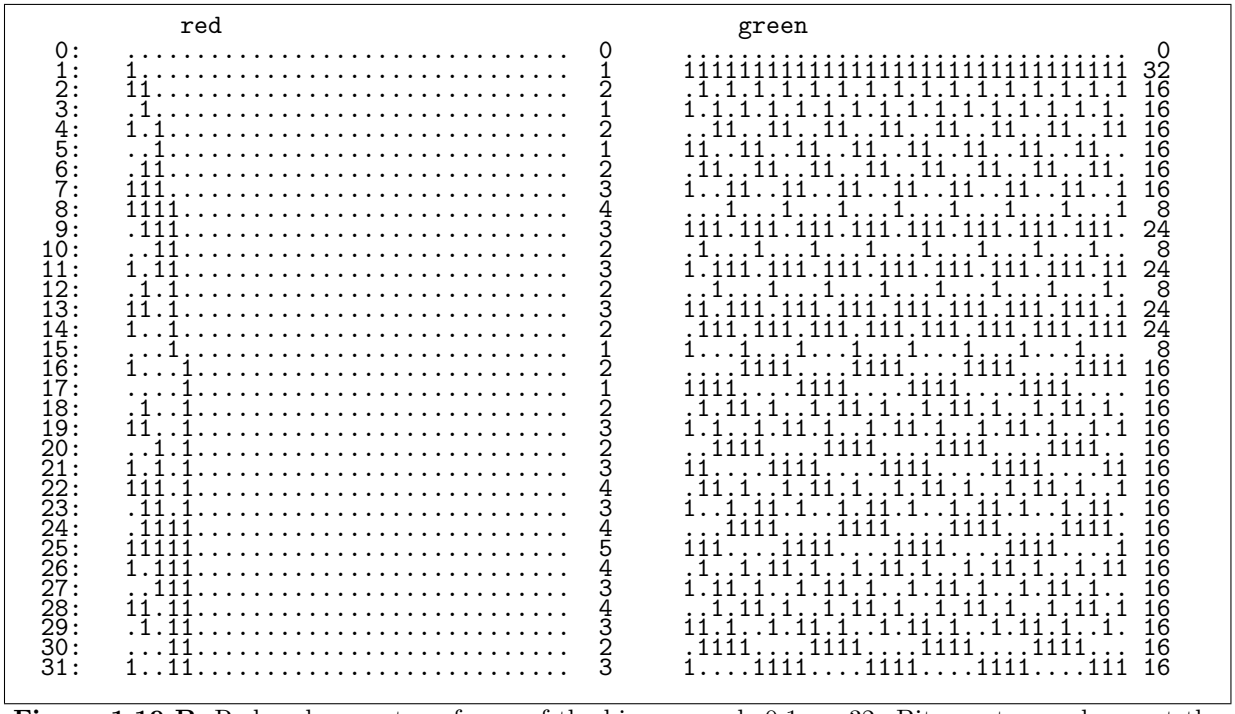

Figure 1.19-B: Red and green transforms of the binary words 0,1,...,32. Bit-counts are shown at the right of each column.

```
1   inline ulong green_code(ulong a)
2   {
3       ulong s = BITS_PER_LONG >> 1;
4       ulong m = ~0UL << s;
5       do
6       {
7           ulong u = a & m;
8           ulong v = a ^ u;
9           a = v ^ (u>>s);
10          a ^= (v<<s);
11          s >>= 1;
12          m ^= (m>>s);
13      }
14      while ( s );
15      return  a;
16  }
```

use the masks

```
........1111111111111111
....1111....1111....1111....1111
..11..11..11..11..11..11..11..11
.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1
```

The transforms of the binary words up to 31 are shown in figure 1.19-B, which was created with the program [FXT: bits/bittransforms-red-demo.cc].

1.19.1 Fixed points of the blue code

The sequence of fixed points of the blue code is (entry A118666 of [240])

```
0, 1, 6, 7, 18, 19, 20, 21, 106, 107, 108, 109, 120, 121, 126, 127, 258, 259, ...
```

If f is a fixed point then f XOR 1 is also a fixed point. Further, $2\,(f \text{ XOR } (2\,f))$ is a fixed point. These facts can be cast into a function that returns a unique fixed point for each argument [FXT: bits/blue-fixed-points.h]:

```
 0 = ......, :    .........., =     0
 1 = .....1, :    .........1, =     1
 2 = ....1., :    .......11., =     6
 3 = ....11, :    .......111, =     7
 4 = ...1.., :    .....1.1.., =    20
 5 = ...1.1, :    .....1..1., =    18
 6 = ...11., :    .....1.1.1, =    21
 7 = ...111, :    .....1..11, =    19
 8 = ..1...., :   ...1111.... =   120
 9 = ..1..1, :    ...11.11.., =   108
10 = ..1.1., :    ...1111111, =   126
11 = ..1.11, :    ...11.1.1., =   106
12 = ..11.., :    ...1111..1, =   121
13 = ..11.1, :    ...11.11.1, =   109
14 = ..111., :    ...1111111, =   127
15 = ..1111, :    ...11.1.11, =   107
16 = .1....., :   .1...1.... =   272
17 = .1...1, :    .1.11.1... =   360
18 = .1..1., :    .1.....1.. =   260
19 = .1..11, :    .1.11111.. =   380
20 = .1.1.., :    .1...1.11. =   278
21 = .1.1.1, :    .1.11.111. =   366
22 = .1.11., :    .1......1. =   258
23 = .1.111, :    .1.1111.1. =   378
24 = .11...., :   .1...1...1 =   273
25 = .11..1, :    .1.11.1..1 =   361
26 = .11.1., :    .1.....1.1 =   261
27 = .11.11, :    .1.11111.1 =   381
28 = .111.., :    .1...1.111 =   279
29 = .111.1, :    .1.11.1111 =   367
30 = .1111., :    .1......11 =   259
31 = .11111, :    .1.1111.11 =   379
```

Figure 1.19-C: The first fixed points of the blue code. The highest bit of all fixed points lies at an even index. There are $2^{n/2}$ fixed points with highest bit at index n.

```
1   inline ulong blue_fixed_point(ulong s)
2   {
3       if ( 0==s )  return 0;
4       ulong f = 1;
5       while ( s>1 )
6       {
7           f ^= (f<<1);
8           f <<= 1;
9           f |= (s&1);
10          s >>= 1;
11      }
12      return f;
13  }
```

The output for the first few arguments is shown in figure 1.19-C. Note that the fixed points are not in ascending order. The list was created by the program [FXT: bits/bittransforms-blue-fp-demo.cc].

Now write $f(x)$ for the binary polynomial corresponding to f (see chapter 38 on page 815), if $f(x)$ is a fixed point (that is, $B\,f(x) = f(x+1) = f(x)$), then both $(x^2+x)\,f(x)$ and $1 + (x^2+x)\,f(x)$ are fixed points. The function `blue_fixed_point()` repeatedly multiplies by x^2+x and adds one if the corresponding bit of the argument is set.

The inverse function uses the fact that polynomial division by $x+1$ can be achieved with the inverse reversed Gray code (see section 1.16.6 on page 45) if the polynomial is divisible by $x+1$:

```
1   inline ulong blue_fixed_point_idx(ulong f)
2   // Inverse of blue_fixed_point()
3   {
4       ulong s = 1;
5       while ( f )
6       {
7           s <<= 1;
8           s ^= (f & 1);
9           f >>= 1;
10          f = inverse_rev_gray_code(f);  // == bitpol_div(f, 3);
11      }
12      return s >> 1;
13  }
```

1.19.2 Relations between the transforms

We write B for the blue code (transform), Y for the yellow code and r for bit-reversal (the `revbin`-function). Then B and Y are connected by the relations

$$B = Y\,r\,Y = r\,Y\,r \tag{1.19-1a}$$
$$Y = B\,r\,B = r\,B\,r \tag{1.19-1b}$$
$$r = Y\,B\,Y = B\,Y\,B \tag{1.19-1c}$$

As said, B and Y are self-inverse:

$$B^{-1} = B, \qquad B\,B = \mathrm{id} \tag{1.19-2a}$$
$$Y^{-1} = Y, \qquad Y\,Y = \mathrm{id} \tag{1.19-2b}$$

The red code and the green code are not involutions ('square roots of identity') but third roots of identity (Using R for the red code, E for the green code):

$$R\,R\,R = \mathrm{id}, \qquad R^{-1} = R\,R = E \tag{1.19-3a}$$
$$E\,E\,E = \mathrm{id}, \qquad E^{-1} = E\,E = R \tag{1.19-3b}$$
$$R\,E = E\,R = \mathrm{id} \tag{1.19-3c}$$

By construction

$$R = r\,B \tag{1.19-4a}$$
$$E = r\,Y \tag{1.19-4b}$$

Relations connecting R and E are:

$$R = E\,r\,E = r\,E\,r \tag{1.19-5a}$$
$$E = R\,r\,R = r\,R\,r \tag{1.19-5b}$$
$$R = R\,E\,R \tag{1.19-5c}$$
$$E = E\,R\,E \tag{1.19-5d}$$

One has

$$r = Y\,R = R\,B = B\,E = E\,Y \tag{1.19-6}$$

Further

$$B = RY = Y\,E = R\,B\,R = E\,B\,E \tag{1.19-7a}$$
$$Y = E\,B = B\,R = R\,Y\,R = E\,Y\,E \tag{1.19-7b}$$

$$R = B\,Y = B\,E\,B = Y\,E\,Y \tag{1.19-7c}$$
$$E = Y\,B = B\,R\,B = Y\,R\,Y \tag{1.19-7d}$$

and

$$\mathrm{id} = B\,Y\,E = R\,Y\,B \tag{1.19-8a}$$
$$\mathrm{id} = E\,B\,Y = B\,R\,Y \tag{1.19-8b}$$
$$\mathrm{id} = Y\,E\,B = Y\,B\,R \tag{1.19-8c}$$

The following multiplication table lists $z = y\,x$. The R in the third column of the second row says that $r\,B = R$. The letter i is used for identity (id). An asterisk says that $x\,y \neq y\,x$.

	i	r	B	Y	R	E
i	i	r	B	Y	R	E
r	r	i	R*	E*	B*	Y*
B	B	E*	i	R*	Y*	r*
Y	Y	R*	E*	i	r*	B*
R	R	Y*	r*	B*	E	i
E	E	B*	Y*	r*	i	R

1.19.3 Relations to Gray code and reversed Gray code

Write g for the Gray code, then:

$$g\,B\,g\,B = \mathrm{id} \quad (1.19\text{-}9a)$$
$$g\,B\,g = B \quad (1.19\text{-}9b)$$
$$g^{-1}\,B\,g^{-1} = B \quad (1.19\text{-}9c)$$
$$g\,B = B\,g^{-1} \quad (1.19\text{-}9d)$$

Let S_k be the operator that rotates a word by k bits (bit zero is moved to position k, use [FXT: `bit_rotate_sgn()` in bits/bitrotate.h]) then

$$Y\,S_{+1}\,Y = g \quad (1.19\text{-}10a)$$
$$Y\,S_{-1}\,Y = g^{-1} \quad (1.19\text{-}10b)$$
$$Y\,S_{k}\,Y = g^{k} \quad (1.19\text{-}10c)$$

Shift in the frequency domain is derivative in time domain. Relation 1.19-10c, together with a algorithm to generate the cycle leaders of the Gray permutation (section 2.8.1 on page 106) gives a curious method to generate the binary necklaces whose length is a power of two, described in section 17.1.6 on page 359. Let e be the operator for the reversed Gray code, then

$$B\,S_{+1}\,B = e^{-1} \quad (1.19\text{-}11a)$$
$$B\,S_{-1}\,B = e \quad (1.19\text{-}11b)$$
$$B\,S_{k}\,B = e^{-k} \quad (1.19\text{-}11c)$$

1.19.4 More transforms by symbolic powering

The idea of powering a transform (as done for the Gray code in section 1.18 on page 48) can be applied to the 'color'-transforms as exemplified for the blue code:

```
1   inline ulong blue_xcode(ulong a, ulong x)
2   {
3       x &= (BITS_PER_LONG-1);  // modulo BITS_PER_LONG
4       ulong s = BITS_PER_LONG >> 1;
5       ulong m = ~0UL << s;
6       while ( s )
7       {
8           if ( x & 1 )  a ^= ( (a&m) >> s );
9           x >>= 1;
10          s >>= 1;
11          m ^= (m>>s);
12      }
13      return  a;
14  }
```

The result is *not* the power of the blue code which would be pretty boring as $B\,B = \mathrm{id}$. Instead the transform (and the equivalents for Y, R and E, see [FXT: bits/bitxtransforms.h]) are more interesting: all relations between the transforms are still valid, if the symbolic exponent is identical with all terms.

For example, we had $B\,B = \mathrm{id}$, now $B^x\,B^x = \mathrm{id}$ is true for all x (there are essentially `BITS_PER_LONG` different x). Similarly, $E\,E = R$ now has to be $E^x\,E^x = R^x$. That is, we have `BITS_PER_LONG` different versions of our four transforms that share their properties with the 'simple' versions. Among them `BITS_PER_LONG` transforms B^x and Y^x that are involutions and E^x and R^x that are third roots of the identity: $E^x\,E^x\,E^x = R^x\,R^x\,R^x = \mathrm{id}$.

While not powers of the simple versions, we still have $B^0 = Y^0 = R^0 = E^0 = \mathrm{id}$. Further, let e be the 'exponent' of all ones and Z be any of the transforms, then $Z^e = Z$, Writing '+' for the XOR operation, then $Z^x\,Z^y = Z^{x+y}$ and so $Z^x\,Z^y = Z$ whenever $x + y = e$.

1.19.5 The building blocks of the transforms

Consider the following transforms on two-bit words where addition is bit-wise (that is, XOR):

$$\mathrm{id}_2\ v = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix} \tag{1.19-12a}$$

$$r_2\,v = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} b \\ a \end{bmatrix} \tag{1.19-12b}$$

$$B_2\,v = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a+b \\ b \end{bmatrix} \tag{1.19-12c}$$

$$Y_2\,v = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ a+b \end{bmatrix} \tag{1.19-12d}$$

$$R_2\,v = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} b \\ a+b \end{bmatrix} \tag{1.19-12e}$$

$$E_2\,v = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a+b \\ a \end{bmatrix} \tag{1.19-12f}$$

It can easily be verified that for these the same relations hold as for id, r, B, Y, R, E. In fact the 'color-transforms', bit-reversion and (trivially) id are the transforms obtained by the repeated Kronecker-products of the matrices. The transforms are linear over GF(2):

$$Z(\alpha\,a + \beta\,b) = \alpha\,Z(a) + \beta\,Z(b) \tag{1.19-13}$$

The corresponding version of the bit-reversal is [FXT: `xrevbin()` in bits/revbin.h]:

```
1   inline ulong xrevbin(ulong a, ulong x)
2   {
3       x &= (BITS_PER_LONG-1);  // modulo BITS_PER_LONG
4       ulong s = BITS_PER_LONG >> 1;
5       ulong m = ~0UL >> s;
6       while ( s )
7       {
8           if ( x & 1 )  a = ( (a & m) << s ) ^ ( (a & (~m)) >> s );
9           x >>= 1;
10          s >>= 1;
11          m ^= (m<<s);
12      }
13      return  a;
14  }
```

Then, for example, $R^x = r^x\,B^x$ (see relation 1.19-4a on page 53). The yellow code is the bit-wise Reed-Muller transform (described in section 22.12 on page 476) of a binary word. The symbolic powering is equivalent to selecting individual levels of the transform.

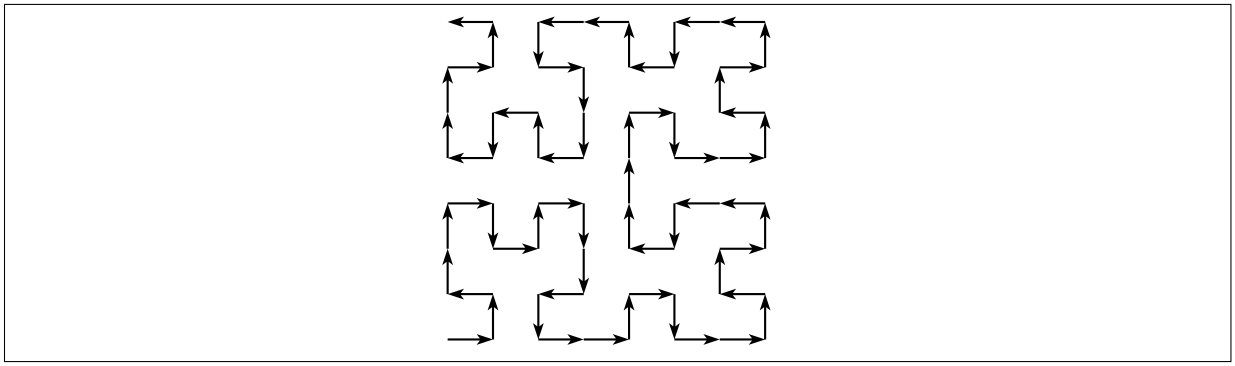

Figure 1.20-A: The Hilbert curve.

```
dx+dy: ++-+++-+++---++++-+++-+++---++++-+++-+++----+---+---+---++++-
dx-dy: +----+++-+++-+++-++++----+---+----++++---+---+----++++---+---+--
  dir: >^<^^>v>^>vv<v>>^>v>>^<^>^<<v<^^^>v>>^<^>^<<v<^<<v>vv<^<v<^^>^<
 turn: 0--+0++--++0+--0-++-0--++--0-++00++-0--++--0-++-0--+0++--++0+--
```

Figure 1.20-B: Moves and turns of the Hilbert curve.

1.20 Moves of the Hilbert curve

A rendering of the Hilbert curve is shown in figure 1.20-A. An efficient algorithm that computes the direction of the n-th move of the Hilbert curve is based on the parity of the number of threes in the radix-4 representation of n (see section 36.9.1 on page 735).

Let d_x and d_y correspond to the moves at step n in the Hilbert curve. Then $d_x, d_y \in \{-1, 0, +1\}$ and exactly one of them is zero. Thereby for both $p := d_x + d_y$ and $m := d_x - d_y$ we have $p, m \in \{-1, +1\}$.

The following function computes p and returns $0, 1$ if $p = -1, +1$, respectively [FXT: bits/hilbert.h]:

```
1   inline ulong hilbert_p(ulong t)
2   // Let dx,dy be the horizontal,vertical move
3   // with step t of the Hilbert curve.
4   // Return  zero if (dx+dy)==-1, else one (then: (dx+dy)==+1).
5   // Algorithm: count number of threes in radix 4
6   {
7       ulong d = (t & 0x5555555555555555UL) & ((t & 0xaaaaaaaaaaaaaaaaUL) >> 1);
8       return  parity( d );
9   }
```

The function can be slightly optimized as follows (64-bit version only):

```
1   inline ulong hilbert_p(ulong t)
2   {
3       t &= ((t & 0xaaaaaaaaaaaaaaaaUL) >> 1);
4       t ^= t>>2;
5       t ^= t>>4;
6       t ^= t>>8;
7       t ^= t>>16;
8       t ^= t>>32;
9       return  t & 1;
10  }
```

The value of m can be computed as:

```
1   inline ulong hilbert_m(ulong t)
2   // Let dx,dy be the horizontal,vertical move
3   // with step t of the Hilbert curve.
4   // Return  zero if (dx-dy)==-1, else one (then: (dx-dy)==+1).
5   {
6       return  hilbert_p( -t );
7   }
```

It remains to merge the values of p and m into a two-bit value d that encodes the direction of the move:

```
1   inline ulong hilbert_dir(ulong t)
2   // Return d encoding the following move with the Hilbert curve.
3   //
4   // d \in {0,1,2,3} as follows:
5   //   d : direction
6   //   0 : right (+x:  dx=+1, dy= 0)
7   //   1 : down  (-y:  dx= 0, dy=-1)
8   //   2 : up    (+y:  dx= 0, dy=+1)
9   //   3 : left  (-x:  dx=-1, dy= 0)
10  {
11      ulong p = hilbert_p(t);
12      ulong m = hilbert_m(t);
13      ulong d = p ^ (m<<1);
14      return  d;
15  }
```

To print the value of d symbolically, one can use the C++ statement `cout << ("v><^")[d];`.

The turn u between steps can be computed as

```
1   inline int hilbert_turn(ulong t)
2   // Return the turn (left or right) with the steps
3   //   t and t-1 of the Hilbert curve.
4   // Returned value is
5   //   0 for no turn
6   //  +1 for right turn
7   //  -1 for left turn
8   {
9       ulong d1 = hilbert_dir(t);
10      ulong d2 = hilbert_dir(t-1);
11      d1 ^= (d1>>1);
12      d2 ^= (d2>>1);
13      ulong u = d1 - d2;
14      // at this point, symbolically:  cout << ("+.-0+.-")[ u + 3 ];
15      if ( 0==u )  return 0;
16      if ( (long)u<0 )  u += 4;
17      return  (1==u ? +1 : -1);
18  }
```

To print the value of u symbolically, one can use `cout << ("-0+")[d+1];`.

The values of p and m, followed by the direction and turn of the Hilbert curve are shown in figure 1.20-B. The list was created with the program [FXT: bits/hilbert-moves-demo.cc]. Figure 1.20-A was created with the program [FXT: bits/hilbert-texpic-demo.cc]. More information about the Hilbert curve can be found in [265, ch.14].

1.21 The Z-order

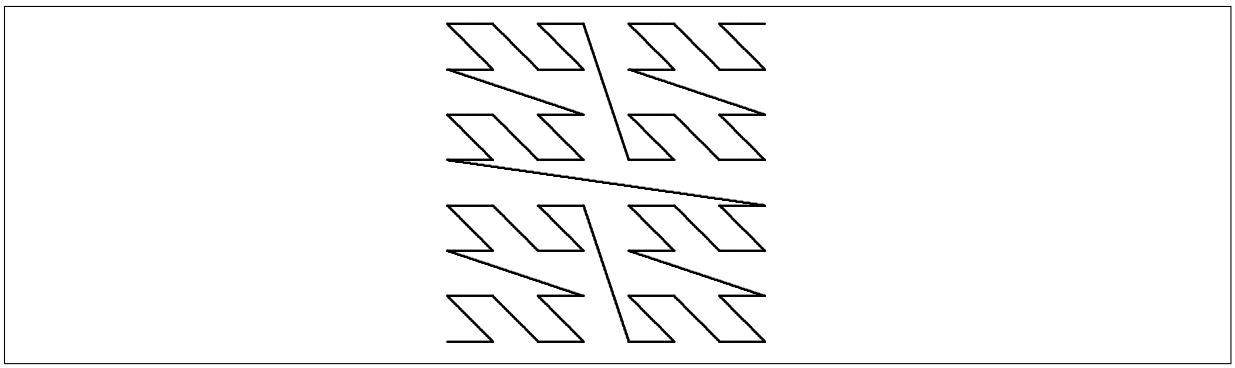

Figure 1.21-A: The Z-order curve.

A 2-dimensional space-filling curve that traverses all points in each quadrant before it enters the next can be obtained by the *Z-order*. Figure 1.21-A shows a rendering of the Z-order curve, it was created

with the program [FXT: bits/zorder-texpic-demo.cc]. The conversion between a linear parameter to a pair coordinates achieved by separating the bits at the even and odd indices [FXT: bits/zorder.h]:

```
inline void lin2zorder(ulong t, ulong &x, ulong &y)  { bit_unzip2(t, x, y); }
```

The routine `bit_unzip2()` is described in section 1.15 on page 39. The inverse is

```
inline ulong zorder2lin(ulong x, ulong y)  { return bit_zip2(x, y); }
```

From any coordinate pair the next pair can be computed with the following (constant amortized time) routine:

```
1   inline void zorder_next(ulong &x, ulong &y)
2   {
3       ulong b = 1;
4       do
5       {
6           x ^= b;  b &= ~x;
7           y ^= b;  b &= ~y;
8           b <<= 1;
9       }
10      while ( b );
11  }
```

The previous pair is obtained similarly:

```
1   inline void zorder_prev(ulong &x, ulong &y)
2   {
3       ulong b = 1;
4       do
5       {
6           x ^= b;  b &= x;
7           y ^= b;  b &= y;
8           b <<= 1;
9       }
10      while ( b );
11  }
```

The routines are written in a way that generalizes trivially to more dimensions:

```
1   inline void zorder3d_next(ulong &x, ulong &y, ulong &z)
2   {
3       ulong b = 1;
4       do
5       {
6           x ^= b;  b &= ~x;
7           y ^= b;  b &= ~y;
8           z ^= b;  b &= ~z;
9           b <<= 1;
10      }
11      while ( b );
12  }
```

```
1   inline void zorder3d_prev(ulong &x, ulong &y, ulong &z)
2   {
3       ulong b = 1;
4       do
5       {
6           x ^= b;  b &= x;
7           y ^= b;  b &= y;
8           z ^= b;  b &= z;
9           b <<= 1;
10      }
11      while ( b );
12  }
```

Unlike the Hilbert curve there are steps where the curve advances more than one unit.

1.22 Scanning for zero bytes

The function (32-bit version)

```
1   static inline ulong contains_zero_byte(ulong x)
2   {
3       return ((x-0x01010101UL)^x) & (~x) & 0x80808080UL;
```

```
4   }
```

from [FXT: bits/zerobyte.h] determines if any sub-byte of the argument is zero. It returns zero when x contains no zero-byte and nonzero when it does. The idea is to subtract one from each of the bytes and then look for bytes where the borrow propagated all the way to the most significant bit. In order to scan for other values than zero (e.g. 0xa5) one can use `contains_zero_byte( x ^ 0xa5a5a5a5UL )`.

Using the simplified version

```
return ((x-0x01010101UL) ^ x) & 0x80808080UL;
```

gives false alarms when a byte equals 0x80. For one byte (in hex, omitting prefixes '0x'):

```
      x-01          = 80-01   = 7f
     (x-01)^x       = 7f ^ 80 = ff
    ((x-01)^x) & 80 = ff & 80 = 80   != 0
```

For strings where the high bit of every byte is known to be zero (for example ASCII-strings) this version can be used. The function [FXT: aux1/bytescan.cc]

```
1   ulong long_strlen(const char *str)
2   // Return length of string starting at str.
3   {
4       ulong x;
5       const char *p = str;
6
7       // Alignment: scan bytes up to word boundary:
8       while ( (ulong)p % BYTES_PER_LONG )
9       {
10          if ( 0 == *p )  return  (ulong)(p-str);
11          ++p;
12      }
13
14      x = *(ulong *)p;
15      while ( ! contains_zero_byte(x) )
16      {
17          p += BYTES_PER_LONG;
18          x = *(ulong *)p;
19      }
20
21      // now a zero byte is somewhere in x:
22      while ( 0 != *p )  { ++p; }
23
24      return  (ulong)(p-str);
25  }
```

may be a win for very long strings and word sizes of 64 or more bits.

1.23 2-adic inverse and square root

1.23.1 Computation of the inverse

The 2-adic inverse can be computed using an iteration (see section 28.1.5 on page 561) with quadratic convergence. The number to be inverted has to be odd [FXT: bits/bit2adic.h]:

```
1   inline ulong inv2adic(ulong x)
2   // Return inverse modulo 2**BITS_PER_LONG
3   // x must be odd
4   // The number of correct bits is doubled with each step
5   // ==> loop is executed prop. log_2(BITS_PER_LONG) times
6   // precision is 3, 6, 12, 24, 48, 96, ... bits (or better)
7   {
8       if ( 0==(x&1) )  return 0;  // not invertible
9       ulong i = x;  // correct to three bits at least
10      ulong p;
11      do
12      {
13          p = i * x;
14          i *= (2UL - p);
15      }
```

```
16      while ( p!=1 );
17      return  i;
18  }
```

Let m be the modulus (a power of two), then the computed value i is the inverse of x modulo m: $i \equiv x^{-1} \bmod m$. It can be used for the so-called *exact division*: to compute the quotient a/x for a number a that is known to be divisible by x, simply multiply by i. This works because $a = b\,x$ (a is divisible by x), so $a\,i \equiv b\,x\,i \equiv b \bmod m$.

1.23.2 Exact division by $C = 2^k \pm 1$

We use the relation (for power series)

$$\frac{A}{C} \;=\; \frac{A}{1-Y} \;=\; A\,(1+Y)\,(1+Y^2)\,(1+Y^4)\,(1+Y^8)\,\ldots\,(1+Y^{2^n}) \mod x^{2^{n+1}} \qquad (1.23\text{-}1)$$

where $Y = 1 - C$. The relation can be used for efficient exact division over $\mathbb{Z}$ by $C = 2^k \pm 1$. For $C = 2^k + 1$ use

$$\frac{A}{C} \;=\; A\,(1-2^k)\,(1+2^{k\,2})\,(1+2^{k\,4})\,(1+2^{k\,8})\,\cdots\,(1+2^{k\,2^u}) \mod 2^N \qquad (1.23\text{-}2)$$

where $k\,2^u \geq N$. For $C = 2^k - 1$ use ($A/C = -A/-C$)

$$\frac{A}{C} \;=\; -A\,(1+2^k)\,(1+2^{k\,2})\,(1+2^{k\,4})\,(1+2^{k\,8})\,\cdots\,(1+2^{k\,2^u}) \mod 2^N \qquad (1.23\text{-}3)$$

The equivalent method for exact division by polynomials (over GF(2)) is given in section 38.1.6 on page 820.

1.23.3 Computation of the square root

```
  x   = ...............................1 =  1
  inv = ...............................1
 sqrt = ...............................1

  x   = 11111111111111111111111111111111 = -1
  inv = 11111111111111111111111111111111

  x   = ............................1. =  2
  x   = 11111111111111111111111111111111. = -2

  x   = ......................i.i.i.i.11 =  3
  inv = i.i.i.i.i.i.i.i.i.i.i.i.i.i.i.i.11

  x   = 11111111111111111111111111111111.1 = -3
  inv = .1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1

  x   = ............................1.. =  4
 sqrt = ............................1.

  x   = 11111111111111111111111111111111.. = -4
```

```
  x   = ..........................1.1 =  5
  inv = ii..ii..ii..ii..ii..ii..ii..ii.1

  x   = 11111111111111111111111111111111.11 = -5
  inv = ..11..11..11..11..11..11..11..11

  x   = ............................11. =  6
  x   = 11111111111111111111111111111111.1. = -6

  x   = ......................i.ii.ii.111 =  7
  inv = i.ii.ii.ii.ii.ii.ii.ii.ii.ii.ii.111

  x   = 11111111111111111111111111111111..1 = -7
  inv = 1..1..1..1..1..1..1..1..1..1..1..1.1
 sqrt = i..iii..i..ii...ii......i.ii.i.1

  x   = ............................1... =  8

  x   = 11111111111111111111111111111111... = -8

  x   = ..........................1..1 =  9
  inv = ..iii...iii...iii...iii...iii..1
 sqrt = iiiiiiiiiiiiiiiiiiiiiiiiiiiiiii.1
```

Figure 1.23-A: Examples of the 2-adic inverse and square root of x where $-9 \leq x \leq +9$. Where no inverse or square root is given, it does not exist.

With the inverse square root we choose the start value to match $\lfloor d/2 \rfloor + 1$ as that guarantees four bits of initial precision. Moreover, we get control to which of the two possible values the inverse square root is finally reached. The argument modulo 8 has to be equal to one.

```
1   inline ulong invsqrt2adic(ulong d)
2   // Return inverse square root modulo 2**BITS_PER_LONG
3   // Must have:  d==1 mod 8
4   // The number of correct bits is doubled with each step
```

```
5   // ==> loop is executed prop. log_2(BITS_PER_LONG) times
6   // precision is 4, 8, 16, 32, 64, ... bits (or better)
7   {
8       if ( 1 != (d&7) )  return 0;  // no inverse sqrt
9       // start value: if d == ****10001 ==> x := ****1001
10      ulong x = (d >> 1) | 1;
11      ulong p, y;
12      do
13      {
14          y = x;
15          p = (3 - d * y * y);
16          x = (y * p) >> 1;
17      }
18      while ( x!=y );
19      return  x;
20  }
```

The square root can be obtained by a final multiplication with d:

```
1   inline ulong sqrt2adic(ulong d)
2   // Return square root modulo 2**BITS_PER_LONG
3   // Must have: d==1 mod 8  or  d==4 mod 32,  d==16 mod 128
4   //   ... d==4**k mod 4**(k+3)
5   // Result undefined if condition does not hold
6   {
7       if ( 0==d )  return  0;
8       ulong s = 0;
9       while ( 0==(d&1) )  { d >>= 1; ++s; }
10      d *= invsqrt2adic(d);
11      d <<= (s>>1);
12      return   d;
13  }
```

Note that the 2-adic square root is something completely different from the integer square root in general. If the argument d is a perfect square then the result equals $\pm\sqrt{d}$. The output of the program [FXT: bits/bit2adic-demo.cc] is shown in figure 1.23-A. For further information on 2-adic (more generally p-adic) numbers see [168], [111], and also [165].

1.24 Radix −2 representation

The radix −2 representation of a number n is

$$n \;=\; \sum_{k=0}^{\infty} t_k\,(-2)^k \tag{1.24-1}$$

where the t_k are zero or one. For integers n the sum is terminating: the highest nonzero t_k is at most two positions beyond the highest bit of the binary representation of the absolute value of n (with two's complement).

1.24.1 Conversion from binary

Item 128 of [31] gives a surprisingly simple algorithm to obtain the coefficients t_k of the radix −2 representation of a binary number [FXT: bits/negbin.h]:

```
1   inline ulong bin2neg(ulong x)
2   // binary --> radix(-2)
3   {
4       const ulong m = 0xaaaaaaaaUL; // 32 bit version
5       x += m;
6       x ^= m;
7       return  x;
8   }
```

An example:

```
 k:     bin(k)    m=bin2neg(k)   g=gray(m)   dec(g)
 0:    ......      ......       ......       0 <= 0
 1:    ......1     ......1      ......1      1 <= 1
 2:    .....1.     ....11.      ....1.1      5
 3:    .....11     ....111      ....1..      4
 4:    ....1..     ....1..      ....11.      2
 5:    ....1.1     ....1.1      ....111      3 <= 5
 6:    ....11.     ..11.1.      ..1.111     19
 7:    ....111     ..11.11      ..1.11.     18
 8:    ...1...     ..11...      ..1.1..     20
 9:    ...1..1     ..11..1      ..1.1.1     21
10:    ...1.1.     ..1111.      ..1...1     17
11:    ...1.11     ..11111      ..1....     16
12:    ...11..     ..111..      ..1..1.     14
13:    ...11.1     ..111.1      ..1..11     15
14:    ...111.     ..1..1.      ..11.11      7
15:    ...1111     ..1..11      ..11.1.      6
16:    ..1....     ..1....      ..11...      8
17:    ..1...1     ..1...1      ..11..1      9
18:    ..1..1.     ..1.11.      ..111.1     13
19:    ..1..11     ..1.111      ..111..     12
20:    ..1.1..     ..1.1..      ..1111.     10
21:    ..1.1.1     ..1.1.1      ..11111     11 <= 21
22:    ..1.11.     11.1.1.     1.11111     75
23:    ..1.111     11.1.11     1.1111.     74
24:    ..11...     11.1...     1.111..     76
25:    ..11..1     11.1..1     1.111.1     77
26:    ..11.1.     11.111.     1.11..1     73
27:    ..11.11     11.1111     1.11...     72
28:    ..111..     11.11..     1.11.1.     70
29:    ..111.1     11.11.1     1.11.11     71
30:    ..1111.     11...1.     1.1..11     79
31:    ..11111     11...11     1.1..1.     78
```

Figure 1.24-A: Radix -2 representations and their Gray codes. Lines ending in '`<=N`' indicate that all values less or equal N occur in the last column up to that point.

```
14 -->  ..1..1. == 16 - 2 == (-2)^4 + (-2)^1
```

The inverse routine is obtained by executing the inverse of the two steps in reversed order:

```
1  inline ulong neg2bin(ulong x)
2  // radix(-2) --> binary
3  // inverse of bin2neg()
4  {
5      const ulong m = 0xaaaaaaaaUL;  // 32-bit version
6      x ^= m;
7      x -= m;
8      return  x;
9  }
```

Figure 1.24-A shows the output of the program [FXT: bits/negbin-demo.cc]. The sequence of Gray codes of the radix -2 representation is a Gray code for the numbers in the range $0, \ldots, k$ for the following values of k (entry A002450 of [240]):

$$k \;=\; 1,\, 5,\, 21,\, 85,\, 341,\, 1365,\, 5461,\, 21845,\, 87381,\, 349525,\, 1398101,\, \ldots,\, (4^n-1)/3$$

1.24.2 Fixed points of the conversion *

```
  0: ..........      64: ....1.....     256: ..1.......     320: ..1.1.....
  1: .........1      65: ....1....1     257: ..1......1     321: ..1.1....1
  4: ........1.      68: ....1...1.     260: ..1.....1.     324: ..1.1...1.
  5: ........1.1     69: ....1...1.1    261: ..1.....1.1    325: ..1.1...1.1
 16: ......1....     80: ....1.1....    272: ..1...1....    336: ..1.1.1....
 17: ......1...1     81: ....1.1...1    273: ..1...1...1    337: ..1.1.1...1
 20: ......1.1..     84: ....1.1.1..    276: ..1...1.1..    340: ..1.1.1.1..
 21: ......1.1.1     85: ....1.1.1.1    277: ..1...1.1.1    341: ..1.1.1.1.1
```

Figure 1.24-B: The fixed points of the conversion and their binary representations (dots denote zeros).

The sequence of fixed points of the conversion starts as

```
0, 1, 4, 5, 16, 17, 20, 21, 64, 65, 68, 69, 80, 81, 84, 85, 256, ...
```

The binary representations have ones only at even positions (see figure 1.24-B). This is the *Moser – De Bruijn sequence*, entry A000695 of [240]. The generating function of the sequence is

$$\frac{1}{1-x}\sum_{j=0}^{\infty}\frac{4^j\,x^{2^j}}{1+x^{2^j}} \quad = \quad x+4\,x^2+5\,x^3+16\,x^4+17\,x^5+20\,x^6+21\,x^7+64\,x^8+65\,x^9+\ldots \tag{1.24-2}$$

The sequence also appears as exponents in the power series

$$\prod_{k=0}^{\infty}\left(1+x^{4^k}\right) \quad = \quad 1+x+x^4+x^5+x^{16}+x^{17}+x^{20}+x^{21}+x^{64}+x^{65}+x^{68}+\ldots \tag{1.24-3}$$

The k-th fixed point can be obtained moving all bits of the binary representation of k to position $2\,x$ where $x \geq 0$ is the index of the bit under consideration:

```
1   inline ulong negbin_fixed_point(ulong k)
2   {
3       return bit_zip0(k);
4   }
```

The used function is given in section 1.15 on page 39. The sequence of radix −2 representations of 0, 1, 2, ..., interpreted as binary numbers, is entry A005351 of [240]:

```
0,1,6,7,4,5,26,27,24,25,30,31,28,29,18,19,16,17,22,23,20,21,106,107,104,105,110,111, ...
```

The corresponding sequence for the negative numbers −1, −2, −3, ... is entry A005352:

```
3,2,13,12,15,14,9,8,11,10,53,52,55,54,49,48,51,50,61,60,63,62,57,56,59,58,37,36,39,38, ...
```

More information about 'non-standard' representations of numbers can be found in [168].

1.24.3 Generating negbin words in order

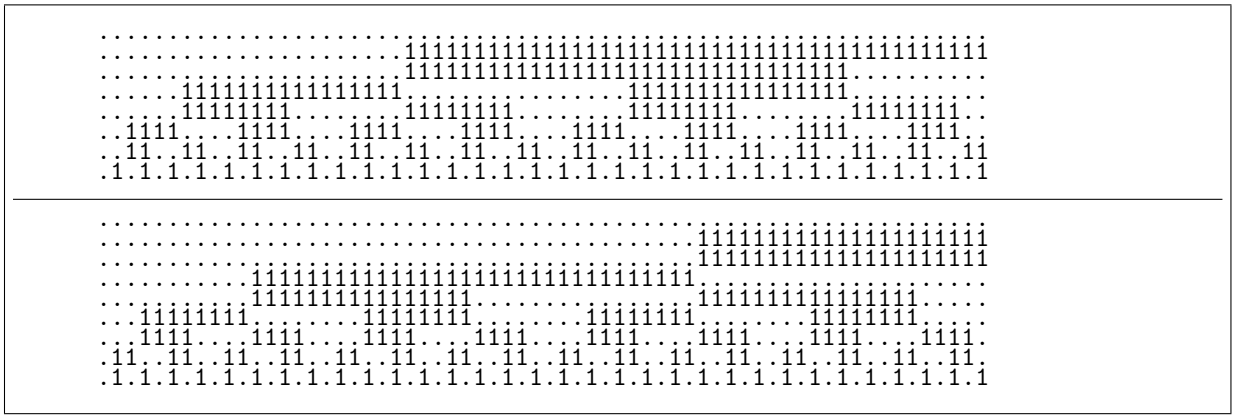

Figure 1.24-C: Radix −2 representations of the numbers $0\ldots+63$ (top), and $0\ldots-63$ (bottom).

A radix −2 representation can be incremented by the function [FXT: bits/negbin.h] (32-bit versions in what follows):

```
1   inline ulong next_negbin(ulong x)
2   // With x the radix(-2) representation of n
3   // return radix(-2) representation of n+1.
4   {
5       const ulong m = 0xaaaaaaaaUL;
6       x ^= m;
7       ++x;
8       x ^= m;
9       return x;
10  }
```

A version without constants is

```
1       ulong s = x << 1;
2       ulong y = x ^ s;
3       y += 1;
4       s ^= y;
5       return s;
```

Decrementing can be done via

```
1   inline ulong prev_negbin(ulong x)
2   // With x the radix(-2) representation of n
3   // return radix(-2) representation of n-1.
4   {
5       const ulong m = 0xaaaaaaaaUL;
6       x ^= m;
7       --x;
8       x ^= m;
9       return x;
10  }
```

or via

```
1       const ulong m = 0x55555555UL;
2       x ^= m;
3       ++x;
4       x ^= m;
5       return x;
```

The functions are quite fast, about 440 million words per second are generated (5 cycles per increment or decrement). Figure 1.24-C shows the generated words in forward (top) and backward (bottom) order, it was created with the program [FXT: bits/negbin2-demo.cc].

1.25 A sparse signed binary representation

```
 0:    .......    .......     0 =
 1:    ......1    ......P     1 =  +1
 2:    .....1.    .....P.     2 =  +2
 3:    .....11    ....P.M     3 =  +4 -1
 4:    ....1..    ....P..     4 =  +4
 5:    ....1.1    ....P.P     5 =  +4 +1
 6:    ....11.    ...P.M.     6 =  +8 -2
 7:    ....111    ...P..M     7 =  +8 -1
 8:    ...1...    ...P...     8 =  +8
 9:    ...1..1    ...P..P     9 =  +8 +1
10:    ...1.1.    ...P.P.    10 =  +8 +2
11:    ...1.11    ..P.M.M    11 =  +16 -4 -1
12:    ...11..    ..P.M..    12 =  +16 -4
13:    ...11.1    ..P.M.P    13 =  +16 -4 +1
14:    ...111.    ..P..M.    14 =  +16 -2
15:    ...1111    ..P...M    15 =  +16 -1
16:    ..1....    ..P....    16 =  +16
17:    ..1...1    ..P...P    17 =  +16 +1
18:    ..1..1.    ..P..P.    18 =  +16 +2
19:    ..1..11    ..P.P.M    19 =  +16 +4 -1
20:    ..1.1..    ..P.P..    20 =  +16 +4
21:    ..1.1.1    ..P.P.P    21 =  +16 +4 +1
22:    ..1.11.    .P.M.M.    22 =  +32 -8 -2
23:    ..1.111    .P.M..M    23 =  +32 -8 -1
24:    ..11...    .P.M...    24 =  +32 -8
25:    ..11..1    .P.M..P    25 =  +32 -8 +1
26:    ..11.1.    .P.M.P.    26 =  +32 -8 +2
27:    ..11.11    .P..M.M    27 =  +32 -4 -1
28:    ..111..    .P..M..    28 =  +32 -4
29:    ..111.1    .P..M.P    29 =  +32 -4 +1
30:    ..1111.    .P...M.    30 =  +32 -2
31:    ..11111    .P....M    31 =  +32 -1
32:    .1.....    .P.....    32 =  +32
```

Figure 1.25-A: Sparse signed binary representations (nonadjacent from, NAF). The symbols '`P`' and '`M`' are respectively used for $+1$ and -1, dots denote zeros.

An algorithm to compute a representation of a number x as

$$x \;=\; \sum_{k=0}^{\infty} s_k \cdot 2^k \quad \text{where} \quad s_k \in \{-1,\, 0,\, +1\} \tag{1.25-1}$$

such that two consecutive digits s_k, s_{k+1} are never simultaneously nonzero is given in [216]. Figure 1.25-A gives the representation of several small numbers, it is the output of [FXT: bits/bin2naf-demo.cc].

We can convert the binary representation of x into a pair of binary numbers that correspond to the positive and negative digits [FXT: bits/bin2naf.h]:

```
1   inline void bin2naf(ulong x, ulong &np, ulong &nm)
2   // Compute (nonadjacent form, NAF) signed binary representation of x:
3   // the unique representation of x as
4   //   x=\sum_{k}{d_k*2^k} where d_j \in {-1,0,+1}
5   //   and no two adjacent digits d_j, d_{j+1} are both nonzero.
6   // np has bits j set where d_j==+1
7   // nm has bits j set where d_j==-1
8   // Thereby:  x = np - nm
9   {
10      ulong xh = x >> 1;  // x/2
11      ulong x3 = x + xh;  // 3*x/2
12      ulong c = xh ^ x3;
13      np = x3 & c;
14      nm = xh & c;
15  }
```

Converting back to binary is trivial:

```
1   inline ulong naf2bin(ulong np, ulong nm)
2   {
3       return ( np - nm );
4   }
```

The representation is one example of a *nonadjacent form* (NAF). A method for the computation of certain nonadjacent forms (w-NAF) is given in [204]. A Gray code for the signed binary words is described in section 12.5 on page 304.

```
         0:    .......    .......     0 =
         1:    .......1   .......P    1 =  +1
         2:    ......1.   ......P.    2 =  +2
         4:    .....1..   .....P..    4 =  +4
         5:    .....1.1   .....P.P    5 =  +4 +1
         8:    ....1...   ....P...    8 =  +8
         9:    ....1..1   ....P..P    9 =  +8 +1
        10:    ....1.1.   ....P.P.   10 =  +8 +2
        16:    ...1....   ...P....   16 =  +16
        17:    ...1...1   ...P...P   17 =  +16 +1
        18:    ...1..1.   ...P..P.   18 =  +16 +2
        20:    ...1.1..   ...P.P..   20 =  +16 +4
        21:    ...1.1.1   ...P.P.P   21 =  +16 +4 +1
        32:    ..1.....   ..P.....   32 =  +32
        33:    ..1....1   ..P....P   33 =  +32 +1
        34:    ..1...1.   ..P...P.   34 =  +32 +2
        36:    ..1..1..   ..P..P..   36 =  +32 +4
        37:    ..1..1.1   ..P..P.P   37 =  +32 +4 +1
        40:    ..1.1...   ..P.P...   40 =  +32 +8
        41:    ..1.1..1   ..P.P..P   41 =  +32 +8 +1
        42:    ..1.1.1.   ..P.P.P.   42 =  +32 +8 +2
        64:    .1......   .P......   64 =  +64
```

Figure 1.25-B: The numbers whose negative part in the NAF representation is zero.

When a binary word contains no consecutive ones then the negative part of the NAF representation is zero. The sequence of values is [0, 1, 2, 4, 5, 8, 9, 10, 16, ...], entry A003714 of [240], see figure 1.25-B. The numbers are called the *Fibbinary numbers*.

1.26 Generating bit combinations

1.26.1 Co-lexicographic (colex) order

```
 word =     set          = set (reversed)
...111 = { 0, 1, 2 }    = { 2 ,1 ,0 }
..1.11 = { 0, 1, 3 }    = { 3 ,1 ,0 }
..11.1 = { 0, 2, 3 }    = { 3 ,2 ,0 }
..111. = { 1, 2, 3 }    = { 3 ,2 ,1 }
.1..11 = { 0, 1, 4 }    = { 4 ,1 ,0 }
.1.1.1 = { 0, 2, 4 }    = { 4 ,2 ,0 }
.1.11. = { 1, 2, 4 }    = { 4 ,2 ,1 }
.11..1 = { 0, 3, 4 }    = { 4 ,3 ,0 }
.11.1. = { 1, 3, 4 }    = { 4 ,3 ,1 }
.111.. = { 2, 3, 4 }    = { 4 ,3 ,2 }
1...11 = { 0, 1, 5 }    = { 5 ,1 ,0 }
1..1.1 = { 0, 2, 5 }    = { 5 ,2 ,0 }
1..11. = { 1, 2, 5 }    = { 5 ,2 ,1 }
1.1..1 = { 0, 3, 5 }    = { 5 ,3 ,0 }
1.1.1. = { 1, 3, 5 }    = { 5 ,3 ,1 }
1.11.. = { 2, 3, 5 }    = { 5 ,3 ,2 }
11...1 = { 0, 4, 5 }    = { 5 ,4 ,0 }
11..1. = { 1, 4, 5 }    = { 5 ,4 ,1 }
11.1.. = { 2, 4, 5 }    = { 5 ,4 ,2 }
111... = { 3, 4, 5 }    = { 5 ,4 ,3 }
```

Figure 1.26-A: Combinations $\binom{6}{3}$ in co-lexicographic order. The reversed sets are sorted.

Given a binary word with k bits set the following routine computes the binary word that is the next combination of k bits in co-lexicographic order. In the co-lexicographic order the reversed sets are sorted, see figure 1.26-A.

The method to determine the successor is to determine the lowest block of ones and move its highest bit one position up. The rest of the block is then moved to the low end of the word [FXT: bits/bitcombcolex.h]:

```
1   inline ulong next_colex_comb(ulong x)
2   {
3       ulong r = x & -x;       // lowest set bit
4       x += r;                 // replace lowest block by a one left to it
5
6       if ( 0==x )  return 0;  // input was last combination
7
8       ulong z = x & -x;       // first zero beyond lowest block
9       z -= r;                 // lowest block  (cf. lowest_block())
10
11      while ( 0==(z&1) )  { z >>= 1; }  // move block to low end of word
12      return  x | (z>>1);     // need one bit less of low block
13  }
```

One could replace the while-loop by a bit scan and shift combination. The combinations $\binom{32}{20}$ are generated at a rate of about 142 million per second. The rate is about 120 M/s for the combinations $\binom{32}{12}$, the rate with $\binom{60}{7}$ is 70 M/s, and with $\binom{60}{53}$ it is 160 M/s.

The following routine computes the predecessor of a combination:

```
1   inline ulong prev_colex_comb(ulong x)
2   // Inverse of next_colex_comb()
3   {
4       x = next_colex_comb( ~x );
5       if ( 0!=x )  x = ~x;
6       return  x;
7   }
```

The first and last combination can be computed via

```
1   inline ulong first_comb(ulong k)
```

```
2   // Return the first combination of (i.e. smallest word with) k bits,
3   // i.e.  00..001111..1 (k low bits set)
4   // Must have:  0 <= k <= BITS_PER_LONG
5   {
6       ulong t = ~0UL >> ( BITS_PER_LONG - k );
7       if ( k==0 )  t = 0;  // shift with BITS_PER_LONG is undefined
8       return t;
9   }
```

and

```
1   inline ulong last_comb(ulong k, ulong n=BITS_PER_LONG)
2   // return the last combination of (biggest n-bit word with) k bits
3   // i.e.  1111..100..00 (k high bits set)
4   // Must have:  0 <= k <= n <= BITS_PER_LONG
5   {
6       return  first_comb(k) << (n - k);
7   }
```

The `if`-statement in `first_comb()` is needed because a shift by more than `BITS_PER_LONG`−1 is undefined by the C-standard, see section 1.1.5 on page 5.

A variant of the presented (colex-) algorithm appears in [31, item 175]. The variant used here avoids the division of the HAKMEM-version and was given by Doug Moore and Glenn Rhoads. The listing in figure 1.26-A can be created with the program [FXT: bits/bitcombcolex-demo.cc]:

```
1       ulong n = 6,  k = 3;
2       ulong last = last_comb(k, n);
3       ulong g = first_comb(k);
4       ulong gg = 0;
5       do
6       {
7           // visit combination given as word g
8           gg = g;
9           g =  next_colex_comb(g);
10      }
11      while ( gg!=last );
```

1.26.2 Lexicographic (lex) order

lex (5, 3)		colex (5, 2)	
word =	set	word =	set
..111	{ 0, 1, 2 }	...11	{ 0, 1 }
.1.11	{ 0, 1, 3 }	..1.1	{ 0, 2 }
1..11	{ 0, 1, 4 }	..11.	{ 1, 2 }
.11.1	{ 0, 2, 3 }	.1..1	{ 0, 3 }
1.1.1	{ 0, 2, 4 }	.1.1.	{ 1, 3 }
11..1	{ 0, 3, 4 }	.11..	{ 2, 3 }
.111.	{ 1, 2, 3 }	1...1	{ 0, 4 }
1.11.	{ 1, 2, 4 }	1..1.	{ 1, 4 }
11.1.	{ 1, 3, 4 }	1.1..	{ 2, 4 }
111..	{ 2, 3, 4 }	11...	{ 3, 4 }

Figure 1.26-B: Combinations $\binom{5}{3}$ in lexicographic order (left columns). The sets are sorted. The binary words corresponding to lexicographic order are obtained from the bit-reversed complements of the words corresponding to the co-lexicographic order of combinations $\binom{5}{2} = \binom{5}{5-3}$.

The binary words corresponding to combinations $\binom{n}{k}$ in lexicographic order can be obtained as the bit-reversed complements of the words for the combinations $\binom{n}{n-k}$ in colex order, see figure 1.26-B. Note a more precise term for the order is subset-lex (for sets written with elements in increasing order). The sequence is identical to the delta-set-colex order backwards.

The program [FXT: bits/bitcomblex-demo.cc] shows how to compute the subset-lex sequence efficiently:

```
1       ulong n = 5,  k = 3;
2       ulong x = first_comb(n-k);        // first colex (n-k choose n)
3       const ulong m = first_comb(n);    // aux mask
```

```
4       const ulong l = last_comb(k, n);  // last colex
5       ulong ct = 0;
6       ulong y;
7       do
8       {
9           y = revbin(~x, n) & m;  // lex order
10          // visit combination given as word y
11          x = next_colex_comb(x);
12      }
13      while ( y != l );
```

The bit-reversal routine `revbin()` is shown in section 1.14 on page 34. Sections 6.1.1 on page 176 and section 6.1.2 give iterative algorithms for combinations (represented by arrays) in lex and colex order, respectively.

1.26.3 Shifts-order

```
 1:  1....      1:  11...      1:  111..      1:  1111.
 2:  .1...      2:  .11..      2:  .111.      2:  .1111
 3:  ..1..      3:  ..11.      3:  ..111      3:  111.1
 4:  ...1.      4:  ...11      4:  11.1.      4:  11.11
 5:  ....1      5:  1.1..      5:  .11.1      5:  1.111
                6:  .1.1.      6:  11..1
                7:  ..1.1      7:  1.11.
                8:  1..1.      8:  .1.11
                9:  .1..1      9:  1.1.1
               10:  1...1     10:  1..11
```

Figure 1.26-C: Combinations $\binom{5}{k}$, for $k = 1, 2, 3, 4$ in shifts-order.

```
 1:  1111...                  18:  .11..11
 2:  .1111..                  19:  11..1.1  < S
 3:  ..1111.                  20:  11...11  < S-2
 4:  ...1111                  21:  1.111..  < S-2
 5:  111.1..  < S             22:  .1.111.
 6:  .111.1.                  23:  ..1.111
 7:  ..111.1                  24:  1.11.1.  < S
 8:  111..1.  < S             25:  .1.11.1
 9:  .111..1                  26:  1.11..1  < S
10:  111...1  < S             27:  1.1.11.  < S-2
11:  11.11..  < S-2           28:  .1.1.11
12:  .11.11.                  29:  1.1.1.1  < S
13:  ..11.11                  30:  1.1..11  < S-2
14:  11.1.1.  < S             31:  1..111.  < S-2
15:  .11.1.1                  32:  .1..111
16:  11.1..1  < S             33:  1..11.1  < S
17:  11..11.  < S-2           34:  1..1.11  < S-2
18:  .11..11                  35:  1...111  < S-2
```

Figure 1.26-D: Updates with combinations $\binom{7}{4}$: simple split '`S`', split second '`S-2`', easy case unmarked.

Figure 1.26-C shows combinations in shifts-order. The order for combinations $\binom{n}{k}$ is obtained from the shifts-order for subsets (section 8.4 on page 208) by discarding all subsets whose number of elements are $\neq k$ and reversing the list order. The first combination is $[1^k 0^{n-k}]$, and the successor is obtained as follows (see figure 1.26-D):

1. Easy case: if the rightmost one is not in position zero (least significant bit) then shift the word to the right and return the combination.
2. Finished?: if the combination is the last one ($[0^n]$, $[0^{n-1}1]$, $[10^{n-k}1^{k-1}]$) then return zero.
3. Shift back: shift the word to the left such that the leftmost one is in the leftmost position (this can be a no-op).
4. Simple split: if the rightmost one is not the least significant bit then move it one position to the right and return the combination.

5. Split second block: move the rightmost bit of the second block (from the right) of ones one position to the right and attach the lowest block of ones and return the combination.

An implementation is given in [FXT: bits/bitcombshifts.h]:

```
class bit_comb_shifts
{
public:
    ulong x_;  // the combination
    ulong s_;  // how far shifted to the right
    ulong n_, k_;  // combinations (n choose k)
    ulong last_;   // last combination

public:
    bit_comb_shifts(ulong n, ulong k)
    {
        n_ = n;  k_ = k;
        first();
    }

    ulong first(ulong n, ulong k)
    {
        s_ = 0;
        x_ =  last_comb(k, n);

        if ( k>1 )  last_ = first_comb(k-1) | (1UL<<(n_-1));  // [10000111]
        else        last_ = k;  //  [000001] or [000000]

        return x_;
    }

    ulong first()  { return first(n_, k_); }

    ulong next()
    {
        if ( 0==(x_&1) )  // easy case:
        {
            ++s_;
            x_ >>= 1;
            return  x_;
        }
        else  // splitting cases:
        {
            if ( x_ == last_ )  return 0;  // combination was last

            x_ <<= s_;  s_ = 0;  // shift back to the left
            ulong b = x_ & -x_;  // lowest bit


            if ( b!=1UL )  // simple split
            {
                x_ -= (b>>1);  // move rightmost bit to the right
                return x_;
            }
            else  // split second block and attach first
            {
                ulong t = low_ones(x_);  //  block of ones at lower end
                x_ ^= t;  // remove block
                ulong b2 = x_ & -x_;  // (second) lowest bit

                b2 >>= 1;
                x_ -= b2;  // move bit to the right

                // attach block:
                do  { t<<=1; }  while ( 0==(t&x_) );
                x_ |= (t>>1);
                return x_;
            }
        }
    }
};
```

The combinations $\binom{32}{20}$ are generated at a rate of about 150 M/s, for the combinations $\binom{32}{12}$ the rate is about 220 M/s [FXT: bits/bitcombshifts-demo.cc]. The rate with $\binom{60}{7}$ is 415 M/s, and with $\binom{60}{53}$ it is 110 M/s.

1.27 Generating bit subsets of a given word

1.27.1 Counting order

In order to generate all bit-subsets of a binary word one can use the sparse counting idea shown in section 1.8 on page 20. The implementation is [FXT: `class bit_subset` in bits/bitsubset.h]:

```
1   class bit_subset
2   {
3   public:
4       ulong u_;  // current subset
5       ulong v_;  // the full set
6
7   public:
8       bit_subset(ulong v) : u_(0), v_(v)  { ; }
9       ~bit_subset() { ; }
10      ulong current()  const  { return u_; }
11      ulong next()  { u_ = (u_ - v_) & v_;  return u_; }
12      ulong prev()  { u_ = (u_ - 1 ) & v_;  return u_; }
13
14      ulong first(ulong v)  { v_=v;  u_=0;  return u_; }
15      ulong first()  { first(v_);  return u_; }
16
17      ulong last(ulong v)  { v_=v;  u_=v;  return u_; }
18      ulong last()  { last(v_);  return u_; }
19  };
```

With the word `[...11.1.]` the following sequence of words is produced by subsequent `next()`-calls:

```
.....1.
....1..
....1.1
...1...
...1..1
...11..
...11.1
........
```

A block of ones at the right will result in the binary counting sequence, see [FXT: bits/bitsubset-demo.cc].

1.27.2 Gray code order

We use a method to isolate the changing bit from counting order that does not depend on shifting:

```
*******0111  =  u
*******1000  =  u+1
00000001111  = (u+1) ^ u
00000001000  = ((u+1) ^ u) & (u+1)  <--= bit to change
```

The method still works if the lowest one are separated by any amount of zeros. In fact, we want to find the single bit that changed from zero to one. The bits that are switched from zero to one in the transition from the word A to B can also be isolated via `X=B&~A`. The implementation is [FXT: `class bit_subset` in bits/bitsubset-gray.h]:

```
1   class bit_subset_gray
2   {
3   public:
4       bit_subset S_;
5       ulong g_;  // subsets in Gray code order
6       ulong h_;  // highest bit in S_.v_;  needed for the prev() method
7
8   public:
9       bit_subset_gray(ulong v) : S_(v), g_(0), h_(highest_one(v))  { ; }
10      ~bit_subset_gray()  { ; }
11
12      ulong current()  const { return g_; }
13      ulong next()
14      {
15          ulong u0 = S_.current();
16          if ( u0 == S_.v_ )  return first();
17          ulong u1 = S_.next();
18          ulong x = ~u0 & u1;
19          g_ ^= x;
```

```
20          return g_;
21      }
22
23      ulong first(ulong v)  { S_.first(v);  h_=highest_one(v);  g_=0;  return g_; }
24      ulong first()  { S_.first();  g_=0;  return g_; }
25    [--snip--]
```

With the word `[...11.1.]` the following sequence of words is produced by subsequent `next()`-calls:

```
....1.
...i1.
...1..
..11..
..11.1
..1..1
..1...
........
```

A block of ones at the right will result in the binary Gray code sequence, see [FXT: bits/bitsubset-gray-demo.cc]. The method `prev()` computes the previous word in the sequence, note the swapped roles of the variables `u0` and `u1`:

```
1     [--snip--]
2       ulong prev()
3       {
4           ulong u1 = S_.current();
5           if ( u1 == 0 )  return last();
6           ulong u0 = S_.prev();
7           ulong x = ~u0 & u1;
8           g_ ^= x;
9           return g_;
10      }
11
12      ulong last(ulong v)  { S_.last(v);  h_=highest_one(v); g_=h_;  return g_; }
13      ulong last()  { S_.last();  g_=h_;  return g_; }
14  };
```

1.28 Binary words as subsets in lexicographic order

1.28.1 Next and previous word in lexicographic order

```
        1:  1...  =  8    {0}
        2:  11..  = 12    {0, 1}
        3:  111.  = 14    {0, 1, 2}
        4:  1111  = 15    {0, 1, 2, 3}
        5:  11.1  = 13    {0, 1, 3}
        6:  1.1.  = 10    {0, 2}
        7:  1.11  = 11    {0, 2, 3}
        8:  1..1  =  9    {0, 3}
        9:  .1..  =  4    {1}
       10:  .11.  =  6    {1, 2}
       11:  .111  =  7    {1, 2, 3}
       12:  .1.1  =  5    {1, 3}
       13:  ..1.  =  2    {2}
       14:  ..11  =  3    {2, 3}
       15:  ...1  =  1    {3}
```

Figure 1.28-A: Binary words corresponding to non-empty subsets of the 4-element set in lexicographic order with respect to subsets. Note the first element of the subsets corresponds to the highest set bit.

The (bit-reversed) binary words in lexicographic order with respect to the subsets can be generated by successive calls to the following function [FXT: bits/bitlex.h]:

```
1   static inline ulong next_lexrev(ulong x)
2   // Return next word in subset-lex order.
3   {
4       ulong x0 = x & -x;  // lowest bit
5       if ( 1!=x0 )  // easy case: set bit right of lowest bit
6       {
```

```
 7          x0 >>= 1;
 8          x ^= x0;
 9          return  x;
10      }
11      else  // lowest bit at word end
12      {
13          x ^= x0;  // delete lowest bit
14          x0 = x & -x;  // new lowest bit ...
15          x0 >>= 1;  x -= x0;  // ... is moved one to the right
16          return  x;
17      }
18  }
```

The bit-reversed representation was chosen because the isolation of the lowest bit is often cheaper than the same operation on the highest bit. Starting with a one-bit word at position $n-1$ one generates the 2^n bit-subsets of length n. The function is used as follows [FXT: bits/bitlex-demo.cc]:

```
ulong n = 4;  // n-bit binary words
ulong x = 1UL<<(n-1);  // first subset
do
{
    // visit word x
}
while ( (x=next_lexrev(x)) );
```

The output for of the program is shown in figure 1.28-A. About 225 million words per second are generated. An equivalent routine for arrays is given in section 8.1.2 on page 202.

```
 [0: ......  =  0 *]   16: .1...1  = 17      32: 1....1  = 33      48: 11..11  = 51
  1: .....1  =  1 *    17: .1..11  = 19      33: 1...11  = 35      49: 11..1.  = 50
  2: ....11  =  3      18: .1..1.  = 18 *    34: 1...1.  = 34 *    50: 11.1.1  = 53
  3: ....1.  =  2      19: .1.1.1  = 21      35: 1..1.1  = 37      51: 11.111  = 55
  4: ...1.1  =  5      20: .1.111  = 23      36: 1..111  = 39      52: 11.11.  = 54
  5: ...111  =  7      21: .1.11.  = 22      37: 1..11.  = 38      53: 11.1..  = 52
  6: ...11.  =  6 *    22: .1.1..  = 20      38: 1..1..  = 36      54: 111..1  = 57
  7: ...1..  =  4      23: .11..1  = 25      39: 1.1..1  = 41      55: 111.11  = 59
  8: ..1..1  =  9      24: .11.11  = 27      40: 1.1.11  = 43      56: 111.1.  = 58
  9: ..1.11  = 11      25: .11.1.  = 26      41: 1.1.1.  = 42      57: 1111.1  = 61
 10: ..1.1.  = 10 *    26: .111.1  = 29      42: 1.11.1  = 45      58: 111111  = 63
 11: ..11.1  = 13      27: .11111  = 31      43: 1.1111  = 47      59: 11111.  = 62
 12: ..1111  = 15      28: .1111.  = 30      44: 1.111.  = 46      60: 1111..  = 60 *
 13: ..111.  = 14      29: .111..  = 28      45: 1.11..  = 44      61: 111...  = 56
 14: ..11..  = 12      30: .11...  = 24      46: 1.1...  = 40      62: 11....  = 48
 15: ..1...  =  8      31: .1....  = 16      47: 11...1  = 49      63: 1.....  = 32
```

Figure 1.28-B: Binary words corresponding to the subsets of the 6-element set, as generated by `prev_lexrev()`. Fixed points are marked with asterisk.

The following function goes backward:

```
1   static inline ulong prev_lexrev(ulong x)
2   // Return previous word in subset-lex order.
3   {
4       ulong x0 = x & -x;  // lowest bit
5       if ( x & (x0<<1) )  // easy case: next higher bit is set
6       {
7           x ^= x0;  // delete lowest bit
8           return x;
9       }
10      else
11      {
12          x += x0;  // move lowest bit to the left
13          x |= 1;   // set rightmost bit
14          return x;
15      }
16  }
```

About 250 million words are generated per second. Starting with zero one obtains, by repeated calls to `prev_lexrev()`, a sequence of words that just before the 2^n-th call has visited every word of length n. The generated sequence of words corresponding to subsets of the 6-element set is shown in figure 1.28-B. The sequence 1, 3, 2, 5, 7, 6, 4, 9,... in the right column is entry A108918 of [240]. It turns out to be useful for a special version of fast Walsh transforms described in section 22.6 on page 462.

1.28.2 Conversion between binary- and lex-ordered words

A little contemplation on the structure of the binary words in lexicographic order leads to the routine that allows random access to the k-th lex-rev word (unrank algorithm) [FXT: bits/bitlex.h]:

```
1   inline ulong negidx2lexrev(ulong k)
2   {
3       ulong z = 0;
4       ulong h = highest_one(k);
5       while ( k )
6       {
7           while ( 0==(h&k) )  h >>= 1;
8           z ^= h;
9           ++k;
10          k &= h - 1;
11      }
12      return  z;
13  }
```

Let the inverse function be $T(x)$, then we have $T(0) = 0$ and, with $h(x)$ being the highest power of two not greater than x,

$$T(x) \quad = \quad h(x) - 1 + \begin{cases} T\,(x - h(x)) & \text{if} \quad x - h(x) \neq 0 \\ h(x) & \text{else} \end{cases} \tag{1.28-1}$$

We obtain the rank algorithm by starting with the lowest bit:

```
1   inline ulong lexrev2negidx(ulong x)
2   {
3       if ( 0==x )  return 0;
4       ulong h = x & -x;  // lowest bit
5       ulong r = (h-1);
6       while ( x^=h )
7       {
8           r += (h-1);
9           h = x & -x;  // next higher bit
10      }
11      r += h;  // highest bit
12      return  r;
13  }
```

1.28.3 Minimal decompositions into terms $2^k - 1$ *

The least number of terms needed in the sum $x = \sum_k \left(2^k - 1\right)$ equals the number of bits of the lex-word as shown in figure 1.28-C. The number can be computed as

```
c = bit_count( negidx2lexrev( x ) );
```

Alternatively, one can subtract the greatest integer of the form $2^k - 1$ until x is zero and count the number of subtractions. The sequence of these numbers is entry A100661 of [240]:

1,2,1,2,3,2,1,2,3,2,3,4,3,2,1,2,3,2,3,4,3,2,3,4,3,4,5,4,3,2,1,2,3,2,3,...

The following function can be used to compute the sequence:

```
1   void S(ulong f, ulong n) // A100661
2   {
3       static int s = 0;
4       ++s;
5       cout << s << ",";
6       for (ulong m=1; m<n; m<<=1)  S(f+m, m);
7       --s;
8       cout << s << ",";
9   }
```

When called with arguments $f = 0$ and $n = 2^k$ it prints the first $2^{k+1} - 1$ numbers of the sequence followed by a zero.

```
...1 1      ...1 =  1 =  1
...11 2     ...1. =  2 =  1 + 1
...1. 1     ...11 =  3 =  3
..1.1 2     ..1.. =  4 =  3 + 1
..111 3     ..1.1 =  5 =  3 + 1 + 1
..11. 2     ..11. =  6 =  3 + 3
..1.. 1     ..111 =  7 =  7
.1..1 2     .1... =  8 =  7 + 1
.1.11 3     .1..1 =  9 =  7 + 1 + 1
.1.1. 2     .1.1. = 10 =  7 + 3
.11.1 3     .1.11 = 11 =  7 + 3 + 1
.1111 4     .11.. = 12 =  7 + 3 + 1 + 1
.111. 3     .11.1 = 13 =  7 + 3 + 3
.11.. 2     .111. = 14 =  7 + 7
.1... 1     .1111 = 15 = 15
1...1 2     1.... = 16 = 15 + 1
1..11 3     1...1 = 17 = 15 + 1 + 1
1..1. 2     1..1. = 18 = 15 + 3
1.1.1 3     1..11 = 19 = 15 + 3 + 1
1.111 4     1.1.. = 20 = 15 + 3 + 1 + 1
1.11. 3     1.1.1 = 21 = 15 + 3 + 3
1.1.. 2     1.11. = 22 = 15 + 7
11..1 3     1.111 = 23 = 15 + 7 + 1
11.11 4     11... = 24 = 15 + 7 + 1 + 1
11.1. 3     11..1 = 25 = 15 + 7 + 3
111.1 4     11.1. = 26 = 15 + 7 + 3 + 1
11111 5     11.11 = 27 = 15 + 7 + 3 + 1 + 1
1111. 4     111.. = 28 = 15 + 7 + 3 + 3
111.. 3     111.1 = 29 = 15 + 7 + 7
11... 2     1111. = 30 = 15 + 15
1.... 1     11111 = 31 = 31
```

Figure 1.28-C: Binary words in subset-lex order and their bit counts (left columns). The least number of terms of the form $2^k - 1$ needed in the sum $x = \sum_k 2^k - 1$ (right columns) equals the bit count.

A generating function of the sequence is given by

$$Z(x) \;:=\; \frac{-1 + 2\,(1-x)\,\prod_{n=1}^{\infty}\left(1 + x^{2^n-1}\right)}{(1-x)^2} \;= \tag{1.28-2}$$
$$1 + 2x + x^2 + 2x^3 + 3x^4 + 2x^5 + x^6 + 2x^7 + 3x^8 + 2x^9 + 3x^{10} + 4x^{11} + 3x^{12} + 2x^{13} + \ldots$$

1.28.4 The sequence of fixed points *

The sequence of fixed points of the conversion to and from indices is 0, 1, 6, 10, 18, 34, 60, 66, 92, 108, 116, 130, 156, 172, 180, 204, 212, 228, 258, 284, 300, 308, 332, 340, 356, 396, 404, 420, 452, 514, 540, 556, ... is sequence A079471 of [240]. Their values as bit patterns are shown in figure 1.28-D. The crucial observation is that a word is a fixed point exactly if (it equals zero or) its bit-count equals 2^j where j is the index of the lowest set bit.

Now we can find out whether x is a fixed point of the sequence by the following function:

```
1   static inline bool is_lexrev_fixed_point(ulong x)
2   // Return whether x is a fixed point in the prev_lexrev() - sequence
3   {
4       if ( x & 1 )
5       {
6           if ( 1==x )  return  true;
7           else         return  false;
8       }
9       else
10      {
11          ulong w = bit_count(x);
12          if ( w != (w & -w) )  return  false;
13          if ( 0==x )  return  true;
14          return  0 != ( (x & -x) & w );
15      }
16  }
```

One can also use either of the following tests:

```
x == negidx2lexrev(x)
x == lexrev2negidx(x)
```

```
          0:  ..........          514:    .1.......1.
          1:  .........1          540:    .1....111..
          6:  ........11.         556:    .1...1.11..
         10:  .......1.1.       [--snip--]
         18:  ......1..1.        1556:   .11....1.1..
         34:  .....1...1.        1572:   .11...1..1..
         60:  .....1111..        1604:   .11..1...1..
         66:  ....1....1.        1668:   .11.1....1..
         92:  ....1.111..        1796:   .111.....1..
        108:  ....11.11..        2040:   .11111111...
        116:  ....111.1..        2050:   1.........1.
        130:  ...1.....1.        2076:   1......111..
        156:  ...1..111..        2092:   1.....1.11..
        172:  ...1.1.11..        2100:   1.....11.1..
        180:  ...1.11.1..        2124:   1....1..11..
        204:  ...11..11..        2132:   1....1.1.1..
        212:  ...11.1.1..        2148:   1....11..1..
        228:  ...111..1..       [--snip--]
        258:  ..1.....1.        4644:   1..1..1..1..
        284:  ..1...111..        4676:  1..1..1...1..
        300:  ..1..1.11..        4740:  1..1.1....1..
        308:  ..1..11.1..        4868:  1..11.....1..
        332:  ..1.1..11..        5112:  1..1111111...
        340:  ..1.1.1.1..        5132:  1.1......11..
        356:  ..1.11..1..        5140:  1.1.....1.1..
        396:  ..11...11..        5156:  1.1....1..1..
        404:  ..11..1.1..        5188:  1.1...1...1..
        420:  ..11.1..1..        5252:  1.1..1....1..
        452:  ..111...1..        5380:  1.1.1.....1..
```

Figure 1.28-D: Fixed points of the binary to lex-rev conversion.

1.28.5 Recursive generation and relation to a power series *

The following function generates the bit-reversed binary words in reversed lexicographic order:

```
1  void C(ulong f, ulong n, ulong w)
2  {
3      for (ulong m=1; m<n; m<<=1)  C(f+m, m, w^m);
4      print_bin("  ", w, 10);  // visit
5  }
```

Calling `C(0, 64, 0)` we obtain the list of words shown in figure 1.28-B with the all-zeros word moved to the last position. A slight modification of the function

```
1  void A(ulong f, ulong n)
2  {
3      cout << "1,";
4      for (ulong m=1; m<n; m<<=1)  A(f+m, m);
5      cout << "0,";
6  }
```

generates the power series (sequence A079559 of [240])

$$\prod_{n=1}^{\infty}\left(1+x^{2^n-1}\right) \quad = \quad 1+x+x^3+x^4+x^7+x^8+x^{10}+x^{11}+x^{15}+x^{16}+\ldots \qquad (1.28\text{-}3)$$

Calling `A(0, 32)` we obtain:

```
1,1,0,1,1,0,0,1,1,0,1,1,0,0,0,1,1,0,1,1,0,0,1,1,0,1,1,0,0,0,0, ...
```

Indeed, the lowest bit of the k-th word of the bit-reversed sequence in reversed lexicographic order equals the $(k-1)$-st coefficient in the power series. The sequence can also be generated with a string substitution engine (see chapter 16 on page 349):

```
Start: 1
Rules:
  0 --> 0
  1 --> 110
-------------
0:   (#=2)
  1
1:   (#=4)
  110
2:   (#=8)
  1101100
```

```
3:   (#=16)
  1101100110110000
4:   (#=32)
  11011001101100011011001101100000
5:   (#=64)
  1101100110110001101100110110000011011001101100011011001101100000
```

We note that the sequence of sums, prepended by one,

$$1 + x\,\frac{\prod_{n=1}^{\infty}\left(1+x^{2^n-1}\right)}{1-x} \quad = \quad 1 + 1\,x + 2\,x^2 + 2\,x^3 + 3\,x^4 + 4\,x^5 + 4\,x^6 + \ldots \tag{1.28-4}$$

has series coefficients

```
1, 1, 2, 2, 3, 4, 4, 4, 5, 6, 6, 7, 8, 8, 8, 8, 9, 10, 10, 11, 12, 12, 12, 13, ...
```

This sequence is entry A046699 of [240]. We have $a(1) = a(2) = 1$ and the sequence satisfies the peculiar recurrence

$$a(n) \quad = \quad a(n - a(n-1)) + a(n - 1 - a(n-2)) \qquad \text{for} \quad n > 2 \tag{1.28-5}$$

1.29 Minimal-change bit combinations

The wonderful routine [FXT: bits/bitcombminchange.h]

```
1   static inline ulong igc_next_minchange_comb(ulong x)
2   // Return the inverse graycode of the next combination in minchange order.
3   // Input must be the inverse graycode of the current combination.
4   {
5       ulong g = rev_gray_code(x);
6       ulong i = 2;
7       ulong cb; // ==candidateBits;
8       do
9       {
10          ulong y = (x & ~(i-1)) + i;
11          ulong j = lowest_one(y) << 1;
12          ulong h = !!(y & j);
13          cb = ((j-h) ^ g) & (j-i);
14          i = j;
15      }
16      while ( 0==cb );
17
18      return  x + lowest_one(cb);
19  }
```

together with

```
1   static inline ulong igc_last_comb(ulong k, ulong n)
2   // Return the (inverse graycode of the) last combination
3   // as in igc_next_minchange_comb()
4   {
5       if ( 0==k )  return 0;
6       ulong f = first_sequency(k);
7       ulong c = first_comb(n);
8       return c ^ (f>>1);
9       // =^=  (by Doug Moore)
10  //    return  ((1UL<<n) - 1) ^ (((1UL<<k) - 1) / 3);
11  }
```

(the function `first_sequency()` is defined in section 1.17 on page 46) can be used as

```
1   static inline ulong next_minchange_comb(ulong x, ulong last)
2   // not efficient, just to explain the usage
3   // of igc_next_minchange_comb()
4   // Must have: last==igc_last_comb(k, n)
5   {
6       x = inverse_gray_code(x);
7       if ( x==last )  return 0;
8       x = igc_next_minchange_comb(x);
9       return  gray_code(x);
10  }
```

Each combination is different from the preceding one in exactly two positions. For example, using $n = 5$ and $k = 3$ we obtain

```
  x       inverse_gray_code(x)
..111       ..1.1 == first_sequency(k)
.11.1       .1..1
.111.       .1.11
.1.11       .11.1
.1111       .1111
11..1       1...1
11.1.       1..11
111..       1.111
1.1.1       11..1
1.11.       11.11
1..11       111.1 == igc_last_comb(k, n)
```

The same run of bit combinations could be obtained by going through the Gray codes and omitting all words where the bit-count is not equal to k. The algorithm shown here, however, is much more efficient.

For reasons of efficiency one may prefer code as

```
1       ulong last = igc_last_comb(k, n);
2       ulong c, nc = first_sequency(k);
3       do
4       {
5           c = nc;
6           nc = igc_next_minchange_comb(c);
7           ulong g = gray_code(c);
8           // Here g contains the bitcombination
9       }
10      while ( c!=last );
```

which avoids the repeated computation of the inverse Gray code.

```
      n = 6  k = 2              n = 6  k = 3              n = 6  k = 4
   ....11  ....1.  ....1.    ...111  ...1.1  ...1..    ..1111  ..1.1.  ..1...
   ...11.  ...1..  ....1.    ..11.1  ..1..1  ....1.    .11.11  .1..1.  ....1.
   ...1.1  ...11.  ....1.    ..111.  ..1.11  ....1.    .1111.  .1.1..  ....1.
   ..11..  ..1...  ...1..    ..1.11  ..11.1  ...1..    .111.1  .1.11.  ...1..
   ..1.1.  ..11..  ....1.    .11..1  .1...1  ...1.    .1.111  .11.1.  ..1...
   ..1..1  ..111.  ....1.    .11.1.  .1..11  ...1..    11..11  1...1.  ....1.
   .11...  .1....  ..1...    .111..  .1.111  ....1.    11.11.  1..1..  ....1.
   .1.1..  .11...  ...1..    .1.1.1  .11..1  ....1.    11.1.1  1..11.  ....1.
   .1..1.  .111..  ....1.    .1.11.  .11.11  ...1.    1111..  1.1...  ...1..
   .1...1  .1111.  ....1.    .1..11  .111.1  ...1..    111.1.  1.11..  ....1.
   11....  1.....  .1....    11...1  1....1  ....1.    111..1  1.111.  ...1..
   1.1...  11....  ..1...    11..1.  1...11  ...1..    1.1.11  11..1.  ....1.
   1..1..  111...  ...1..    11.1..  1..111  ..1...    1.111.  11.1..  ....1.
   1...1.  1111..  ....1.    111...  1.1111  ....1.    1.11.1  11.11.  ...1..
   1....1  11111.  ....1.    1.1..1  11...1  ...1.    1..111  111.1.  ..1...
                             1.1.1.  11..11  ...1..
                             1.11..  11.111  ....1.
                             1..1.1  111..1  ....1.
                             1..11.  111.11  ...1.
                             1...11  1111.1  ...1..
```

Figure 1.29-A: Minimal-change combinations, their inverse Gray codes and the differences of the inverse Gray codes. The differences are powers of two.

The algorithm in `igc_next_minchange_comb()` uses the fact that the difference of two (inverse Gray codes of) successive combinations is always a power of two, see figure 1.29-A. See also [FXT: bits/bitcombminchange-demo.cc]. Using this observation one can derive a different version that checks the pattern of the change:

```
1   static inline ulong igc_next_minchange_comb(ulong x)
2   // Alternative version.
3   {
4       ulong gx = gray_code( x );
5       ulong i = 2;
6       do
7       {
8           ulong y = x + i;
9           i <<= 1;
10          ulong gy = gray_code( y );
11          ulong r = gx ^ gy;
12
13          // Check that change consists of exactly one bit
14          // of the new and one bit of the old pattern:
```

```
15          if ( is_pow_of_2( r & gy ) && is_pow_of_2( r & gx ) )  break;
16          // is_pow_of_2(x):=((x & -x) == x)  returns 1 also for x==0.
17          // But this cannot happen for both tests at the same time
18      }
19      while ( 1 );
20      return  y;
21  }
```

This version is the fastest, the combinations $\binom{32}{12}$ are generated at a rate of about 96 million per second, the combinations $\binom{32}{20}$ at a rate of about 83 million per second.

Here is another version which needs the number of set bits as a second parameter:

```
1   static inline ulong igc_next_minchange_comb(ulong x, ulong k)
2   // Alternative version, uses the fact that the difference
3   // of two successive x is the smallest possible power of 2.
4   {
5       ulong y, i = 2;
6       do
7       {
8           y = x + i;
9           i <<= 1;
10      }
11      while ( bit_count( gray_code(y) ) != k );
12      return  y;
13  }
```

The routine will be fast if the CPU used has a bitcount instruction. The necessary modification for the generation of the previous combination is trivial:

```
1   static inline ulong igc_prev_minchange_comb(ulong x, ulong k)
2   // Returns the inverse graycode of the previous combination in minchange order.
3   // Input must be the inverse graycode of the current combination.
4   // With input==first the output is the last for n=BITS_PER_LONG
5   {
6       ulong y, i = 2;
7       do
8       {
9           y = x - i;
10          i <<= 1;
11      }
12      while ( bit_count( gray_code(y) ) != k );
13      return  y;
14  }
```

1.30 Fibonacci words

A Fibonacci word is a word that does not contain two successive ones. Whether a given binary word is a Fibonacci word can be tested with the function [FXT: bits/fibrep.h]

```
1   inline bool is_fibrep(ulong f)
2   {
3       return  ( 0==(f&(f>>1)) );
4   }
```

1.30.1 Lexicographic order

To generate all Fibonacci words use the following functions from [FXT: bits/fibrep.h]. For forward order (see figure 1.30-A):

```
1   inline ulong next_fibrep(ulong x)
2   // With x the Fibonacci representation of n
3   // return Fibonacci representation of n+1.
4   {
5       // 2 examples:         //  ex. 1             //  ex.2
6       //                     // x == [*]0 010101   // x == [*]0 01010
7       ulong y = x | (x>>1);  // y == [*]? 011111   // y == [*]? 01111
8       ulong z = y + 1;       // z == [*]? 100000   // z == [*]? 10000
```

```
    0: ........     11: ...1.1..     22: .1.....1     33: .1.1.1.1     44: 1..1..1.
    1: .......1     12: ...1.1.1     23: .1....1.     34: 1.......     45: 1..1.1..
    2: ......1.     13: ..1.....     24: .1...1..     35: 1......1     46: 1..1.1.1
    3: .....1..     14: ..1....1     25: .1..,1.1     36: 1.....1.     47: 1.1.....
    4: .....1.1     15: ..1...1.     26: .1..1...     37: 1....1..     48: 1.1....1
    5: ....1...     16: ..1..1..     27: .1..1..1     38: 1....1.1     49: 1.1...1.
    6: ....1..1     17: ..1..1.1     28: .1..1.1.     39: 1...1...     50: 1.1..1..
    7: ....1.1.     18: ..1.1...     29: .1.1....     40: 1...1..1     51: 1.1..1.1
    8: ...1....     19: ..1.1..1     30: .1.1...1     41: 1...1.1.     52: 1.1.1...
    9: ...1...1     20: ..1.1.1.     31: .1.1..1.     42: 1..1....     53: 1.1.1..1
   10: ...1..1.     21: .1......     32: .1.1.1..     43: 1..1...1     54: 1.1.1.1.
```

Figure 1.30-A: All 55 Fibonacci words with 8 bits in lexicographic order.

```
9       z = z & -z;             // z == [0]0 100000   // z == [0]0 10000
10      x ^= z;                 // x == [*]0 110101   // x == [*]0 11010
11      x &= ~(z-1);            // x == [*]0 100000   // x == [*]0 10000
12
13      return x;
14  }
```

The routine can be used as shown in [FXT: bits/fibrep2-demo.cc]:

```
ulong n = 7;
const ulong f = 1UL << n;
ulong t = 0;
do
{
    // visit(t)
    t = next_fibrep(t);
}
while ( t!=f );
```

The reversed order can be obtained via

```
ulong f = 1UL << n;
while ( f )
{
    ulong t = prev_fibrep(f);
    f = t;
    // visit(t)
}
```

which uses the function (64-bit version)

```
1   inline ulong prev_fibrep(ulong x)
2   // With x the Fibonacci representation of n
3   // return Fibonacci representation of n-1.
4   {
5       ulong i = lowest_one_idx(x);
6       x ^= (1UL<<i);
7       ++i;
8       x ^= (0x5555555555555555UL >> (BITS_PER_LONG-i));
9       return x;
10  }
```

The forward version generates about 200 million words per second, the backward version about 135 million words per second.

1.30.2 Gray code order

A Gray code for the binary Fibonacci words (shown in figure 1.30-B) can be obtained by using the Gray code of the radix -2 representations (see section 1.24 on page 61) of binary words whose difference is of the form

```
    1    ...............1
    3    ..............,11
    5    .............1.1
    9    ............1..1
   19    ...........1..11
   37    ..........1..1.1
   73    .........1..1..1
  147    ........1..1..11
  293    .......1..1..1.1
  585    ......1..1..1..1
 1171    .....1..1..1..11
 2341    ....1..1..1..1.1
```

```
 j:        k(j)      k(j)-k(j-1)   x=bin2neg(k)    gray(x)
 1:    ....11...1    .........     ...111...1    ...1..1..1 =  27
 2:    ....11....    ........1     ...111....    ...1..1... =  26
 3:    ....1.1111    ........1     ...111..11    ...1..1.1. =  28
 4:    ....1.11..    .......11     ...11111..    ...1....1. =  23
 5:    ....1.1.11    ........1     ...1111111    ...1..... =  21
 6:    ....1.1.1.    ........1     ...111111.    ...1.....1 =  22
 7:    ....1.1..1    ........1     ...1111..1    ...1...1.1 =  25
 8:    ....1.1...    ........1     ...1111...    ...1...1.. =  24
 9:    ....1...11    .......1.1    ...11..111    ...1.1.1.. =  32
10:    ....1...1.    ........1     ...11..11.    ...1.1.1.1 =  33
11:    ....1....1    ........1     ...11....1    ...1.1...1 =  30
12:    ....1.....    ........1     ...11.....    ...1.1.... =  29
13:    .....11111    ........1     ...11...11    ...1.1..1. =  31
14:    ......11..    .....1..11    .....111..    .....1..1. =  10
15:    ......1.11    ........1     .....11111    .....1.... =   8
16:    ......1.1.    ........1     .....1111.    .....1...1 =   9
17:    ......1..1    ........1     .....11..1    .....1.1.1 =  12
18:    ......1...    ........1     .....11...    .....1.1.. =  11
19:    ........11    .......1.1    .......111    .......1.. =   3
20:    ........1.    ........1     .......11.    .......1.1 =   4
21:    .........1    ........1     .........1    .........1 =   1
22:    ..........    ........1     ..........    .......... =   0
23:    1111111111    ........1     .........11   .........1. =   2
24:    11111111..    ........11    ......11..    ......1.1. =   7
25:    1111111.11    ........1     ......1111    ......1... =   5
26:    1111111.1.    ........1     ......111.    ......1..1 =   6
27:    111111...1    ......1..1    ....11...1    ....1.1..1 =  19
28:    111111....    ........1     ....11....    ....1.1... =  18
29:    11111.1111    ........1     ....11..11    ....1.1.1. =  20
30:    11111.11..    ........11    ....1111..    ....1...1. =  15
31:    11111.1.11    ........1     ....111111    ....1..... =  13
32:    11111.1.1.    ........1     ....11111.    ....1....1 =  14
33:    11111.1..1    ........1     ....111..1    ....1..1.1 =  17
34:    11111.1...    ........1     ....111...    ....1..1.. =  16
```

Figure 1.30-B: Gray code for the binary Fibonacci words (rightmost column).

```
4681    ...1..1..1..1..1
9363    ..1..1..1..1..11
```

The algorithm is to try these values as increments starting from the least, same as for the minimal-change combination described in section 1.29 on page 76. The next valid word is encountered if it is a valid Fibonacci word, that is, if it does not contain two consecutive set bits. The implementation is [FXT: `class bit_fibgray` in bits/bitfibgray.h]:

```
1   class bit_fibgray
2   // Fibonacci Gray code with binary words.
3   {
4   public:
5       ulong x_;  // current Fibonacci word
6       ulong k_;  // aux
7       ulong fw_, lw_;  // first and last Fibonacci word in Gray code
8       ulong mw_;  // max(fw_, lw_)
9       ulong n_;  // Number of bits
10
11  public:
12      bit_fibgray(ulong n)
13      {
14          n_ = n;
15
16          fw_=0;
17          for (ulong m=(1UL<<(n-1)); m!=0; m>>=3)  fw_ |= m;
18          lw_ = fw_ >> 1;
19          if ( 0==(n&1) )  { ulong t=fw_; fw_=lw_; lw_=t; }
20          mw_ = ( lw_>fw_ ? lw_ : fw_ );
21          x_ = fw_;
22
23          k_ = inverse_gray_code(fw_);
24          k_ = neg2bin(k_);
25      }
26
27      ~bit_fibgray() {;}
28
29      ulong next()
30      // Return next word in Gray code.
31      // Return ~0 if current word is the last one.
```

```
32      {
33          if ( x_ == lw_ )  return ~0UL;
34
35          ulong s = n_;
36          while ( 1 )
37          {
38              --s;
39              ulong c = 1 | (mw_ >> s);
40              ulong i = k_ - c;
41              ulong x = bin2neg(i);
42              x ^= (x>>1);
43
44              if ( 0==(x&(x>>1)) )
45              {
46                  k_ = i;
47                  x_ = x;
48                  return x;
49              }
50          }
51      }
52  };
```

More than 100 million words per second are generated. The program [FXT: bits/bitfibgray-demo.cc] shows how to use the class, figure 1.30-B was created with it. Section 12.1 on page 296 gives a recursive algorithm for Fibonacci words in Gray code order.

1.31 Binary words and parentheses strings *

```
 0    .... P   [empty string]        .....   [empty string]
 1    ...1 P   ()                    ....1   ()
 2    ..1.                           ...11   (())
 3    ..11 P   (())                  ..1.1   ()()
 4    .1..                           ..111   ((()))
 5    .1.1 P   ()()                  .1.11   (()())
 6    .11.                           .11.1   ()(())
 7    .111 P   ((()))                .1111   (((())))
 8    1...                           1..11   (())()
 9    1..1                           1.1.1   ()()()
10    1.1.                           1.111   ((()()))
11    1.11 P   (()())                11.11   (()(()))
12    11..                           111.1   ()((()))
13    11.1 P   ()(())                11111   ((((()))))
14    111.
15    1111 P   ((()))
```

Figure 1.31-A: Left: some of the 4-bit binary words can be interpreted as a string parentheses (marked with 'P'). Right: all 5-bit words that correspond to well-formed parentheses strings.

A subset of the binary words can be interpreted as a (well formed) string of parentheses. The 4-bit binary words that have this property are marked with a 'P' in figure 1.31-A (left) [FXT: bits/parenword-demo.cc]. The strings are constructed by scanning the word from the low end and printing a '(' with each one and a ')' with each zero. In order to find out when to terminate one adds up +1 for each opening parenthesis and −1 for a closing parenthesis. When the ones in the binary word have been scanned then s closing parentheses have to be added where s is the value of the sum [FXT: bits/parenwords.h]:

```
1   inline void parenword2str(ulong x, char *str)
2   {
3       int s = 0;
4       ulong j = 0;
5       for (j=0; x!=0; ++j)
6       {
7           s += ( x&1 ? +1 : -1 );
8           str[j] = ")("[x&1];
9           x >>= 1;
10      }
11      while ( s-- > 0 )  str[j++] = ')';  // finish string
12      str[j] = 0;  // terminate string
13  }
```

The 5-bit binary words that are valid 'paren words' together with the corresponding strings are shown in figure 1.31-A (right). Note that the lower bits in the word (right end) correspond to the beginning of the string (left end). If a negative value for the sums occurs at any time of the computation then the word is not a paren word. We use that fact to create a routine that determines whether a word is a paren word:

```
1   inline bool is_parenword(ulong x)
2   {
3       int s = 0;
4       for (ulong j=0; x!=0; ++j)
5       {
6           s += ( x&1 ? +1 : -1 );
7           if ( s<0 )  break;  // invalid word
8           x >>= 1;
9       }
10      return  (s>=0);
11  }
```

The sequence

```
1, 3, 5, 7, 11, 13, 15, 19, 21, 23, 27, 29, 31, 39, 43, 45, 47, 51, 53, 55, 59, 61, 63, ...
```

of nonzero integers x so that `is_parenword(x)` returns `true` is entry A036991 of [240]. If we fix the number of paren pairs then the following functions generate the least and biggest valid binary word. The function simply returns a word with a block of n ones at the low end:

```
1   inline ulong first_parenword(ulong n)
2   // Return least binary word corresponding to n pairs of parens
3   // Example, n=5:  .....11111   (((((())))))
4   {
5       return first_comb(n);
6   }
```

The last paren word is the word with a sequence of n blocks '01' at the low end:

```
1   inline ulong last_parenword(ulong n)
2   // Return biggest binary word corresponding to n pairs of parens.
3   // Must have: 1 <= n <= BITS_PER_LONG/2.
4   // Example, n=5:  .1.1.1.1.1   ()()()()()
5   {
6       return  0x5555555555555555UL >> (BITS_PER_LONG-2*n);
7   }
```

```
  ......11111 = (((((())))))    ...1...1111 = ((((()))()))    ..1....1111 = (((((())))()
  .....1.1111 = ((((()))))    ...1..1.111 = (((()))()))    ..1...1.111 = (((()))))()
  .....11.111 = (((()))))    ...1..11.11 = ((()))()))    ..1...11.11 = ((()))))()
  .....111.11 = ((())))))    ...1..111.1 = ()((()))())    ..1...111.1 = ()((()))))()
  .....1111.1 = ()(((())))    ...1.1..111 = (((())()))    ..1..1..111 = (((()))())()
  ....1..1111 = ((((()))))    ...1.1.1.11 = ((()()()))    ..1..1.1.11 = ((()()))()
  ....1.1.111 = (((()())))    ...1.1.11.1 = ()((()()))    ..1..1.11.1 = ()((()))()
  ....1.11.11 = ((())))))    ...1.11..11 = (())(()))    ..1..11..11 = (())((()))
  ....1.111.1 = ()(((())))    ...1.11.1.1 = ()()(()))    ..1..11.1.1 = ()()((()))
  ....11..111 = (((())))))    ...11...111 = ((()))(()))    ..1.1...111 = ((()))()()
  ....11.1.11 = ((()))))))    ...11..1.11 = ((())))))    ..1.1..1.11 = ((()))()()
  ....11.11.1 = ()((()))))    ...11..11.1 = ()((()))))    ..1.1..11.1 = ()((()))()()
  ....111..11 = (())((())))    ...11.1..11 = (())()(()))    ..1.1.1..11 = (())()(())
  ....111.1.1 = ()()((()))    ...11.1.1.1 = ()()()(()))    ..1.1.1.1.1 = ()()()()()
```

Figure 1.31-B: The 42 binary words corresponding to all valid pairings of 5 parentheses, in colex order.

The sequence of all binary words corresponding to n pairs of parens in colex order can be generated with the following (slightly cryptic) function:

```
1   inline ulong next_parenword(ulong x)
2   // Next (colex order) binary word that is a paren word.
3   {
4       if ( x & 2 )  // Easy case, move highest bit of lowest block to the left:
5       {
6           ulong b = lowest_zero(x);
7           x ^= b;
8           x ^= (b>>1);
9           return x;
10      }
11      else // Gather all low "01"s and split lowest nontrivial block:
```

```
12      {
13          if ( 0==(x & (x>>1)) )  return 0;
14          ulong w = 0;  // word where the bits are assembled
15          ulong s = 0;  // shift for lowest block
16          ulong i = 1;  // == lowest_one(x)
17          do  // collect low "01"s:
18          {
19              x ^= i;
20              w <<= 1;
21              w |= 1;
22              ++s;
23              i <<= 2;  // == lowest_one(x);
24          }
25          while ( 0==(x&(i<<1)) );
26
27          ulong z = x ^ (x+i);  // lowest block
28          x ^= z;
29          z &= (z>>1);
30          z &= (z>>1);
31          w ^= (z>>s);
32          x |= w;
33          return x;
34      }
35  }
```

The program [FXT: bits/parenword-colex-demo.cc] shows how to create a list of binary words corresponding to n pairs of parens (code slightly shortened):

```
1       ulong n = 4;  // Number of paren pairs
2       ulong pn = 2*n+1;
3       char *str = new char[n+1];  str[n] = 0;
4       ulong x = first_parenword(n);
5       while ( x )
6       {
7           print_bin("  ", x, pn);
8           parenword2str(x, str);
9           cout << " = " << str << endl;
10
11          x = next_parenword(x);
12      }
```

Its output with $n = 5$ is shown in figure 1.31-B. The $1,767,263,190$ paren words for $n = 19$ are generated at a rate of about 153 million words per second. Chapter 13 on page 313 gives a different formulation of the algorithm.

Knuth gives [170] a very elegant routine for generating the next paren word, the comments are `MMIX` instructions:

```
1   inline ulong next_parenword(ulong x)
2   {
3       const ulong m0 = -1UL/3;
4       ulong t = x ^ m0;               // XOR t, x, m0;
5       if ( (t&x)==0 )  return 0;      // current is last
6       ulong u = (t-1) ^ t;            // SUBU u, t, 1;  XOR u, t, u;
7       ulong v = x | u;                // OR v, x, u;
8       ulong y = bit_count( u & m0 );  // SADD y, u, m0;
9       ulong w = v + 1;                // ADDU w, v, 1;
10      t = v & ~w;                     // ANDN t, v, w;
11      y = t >> y;                     // SRU y, t, y;
12      y += w;                         // ADDU y, w, y;
13      return y;
14  }
```

The routine is slower, however, about 81 million words per second are generated. A bit-count instruction in hardware would speed it up significantly. By treating the case of easy update separately as in the other version, a rate of about 137 million words per second is obtained.

1.32 Error detection by hashing: the CRC

A *hash value* is an element from a set H that is computed via a *hash function* f that maps any (finite) sequence of input data to H.

For the sake of simplicity we now consider the case that the input sequences are of fixed size so they are in a fixed set, say S. We further assume that the set S is (much) bigger than H.

$$f : S \to H, \qquad s \mapsto h \tag{1.32-1}$$

where $s \in S$ and $h \in H$.

Two input sequences with different hash values are necessarily different. But, as the hash function maps a larger set to a smaller one, there are different input sequences with identical hash values.

A trivial example is the set $H = \{0, 1\}$ together with a function that count binary digits modulo two, the parity function. Another example is the sum-of-digits test (see section 27.5 on page 554) used to check the multiplication of large numbers. In the test we compute the value of a multi digit number decimal $s \in S$ modulo nine. The crucial additional property of this hash is that with $f(A) = a$, $f(B) = b$, $f(C) = c$ (were A, B, C are decimal numbers), then $A \cdot B = C$ implies $a \cdot b = c$.

A hash function f that is actually useful should have the *mixing* property: it should map the elements $s \in S$ 'randomly' to H. With the sum-of-digits test we could have used rather arbitrary moduli for the hash function. With one exception: the value modulo ten as hash would be pretty useless. No change in any digit except for the last could ever be detected.

The so-called *cyclic redundancy check* (CRC) is a hash where the hash values are binary words of fixed length. The hash function (basically) computes $h = s \bmod c$ where s is a binary polynomial build from the input sequence and c is a binary polynomial that is primitive (see chapter 38 on page 815). We will use polynomials c of degree 64 so the hash values (CRCs) are 64-bit words.

An C++ implementation is given as [FXT: `class crc64` in bits/crc64.h]:

```
1   class crc64
2   // 64-bit CRC (cyclic redundancy check)
3   {
4   public:
5       uint64 a_;  // internal state (polynomial modulo c)
6       uint64 c_;  // a binary primitive polynomial
7       // (non-primitive c lead to smaller periods)
8       // The leading coefficient needs not be present.
9       uint64 h_;  // auxiliary
10
11      static const uint64 cc[];  // 16 "random" 64-bit primitive polynomials
12
13  public:
14      crc64(uint64 c=0)
15      {
16          if ( 0==c )  c = 0x1bULL; // =^= 64,4,3,1,0 (default)
17          init(c);
18      }
19
20      ~crc64()  {;}
21
22      void init(uint64 c)
23      {
24          c_ = c;
25          c_ >>= 1;
26          h_ = 1ULL<<63;
27          c_ |= h_;  // leading coefficient
28          reset();
29      }
30
31      void reset()  { set_a(~0ULL); }  // all ones
32      void set_a(uint64 a)  { a_=a; }
33
34      uint64 get_a()  const  { return a_; }
35     [--snip--]
```

Note that a nonzero initial state (member variable `a`) is used: starting with zero will only go to a nonzero state with the first nonzero bit in the input sequence. That is, input sequences differing only by initial runs of zeros would get the same CRC.

After an instance is created one can feed in bits via `bit_in(b)` where lowest bit of `b` must contain the bit to be used:

```
1       [--snip--]
2        void shift()
3        {
4            bool s = (a_ & 1);
5            a_ >>= 1;
6            if ( 0!=s )  a_ ^= c_;
7        }
8
9        uint64 bit_in(unsigned char b)
10       {
11           a_ ^= (b&1);
12           shift();
13           return  a_;
14       }
15      [--snip--]
```

When a byte is to be checksummed we can do better than just feeding in the bits one by one. This is achieved adding the byte followed by eight calls to `shift()`:

```
1       [--snip--]
2        uint64 byte_in(unsigned char b)
3        {
4   #if 1
5            a_ ^= b;
6            shift();  shift();  shift();  shift();
7            shift();  shift();  shift();  shift();
8   #else  // identical but slower:
9            bit_in(b);  b>>=1;  // bit 0
10           bit_in(b);  b>>=1;  // bit 1
11           bit_in(b);  b>>=1;  // bit 2
12           bit_in(b);  b>>=1;  // bit 3
13           bit_in(b);  b>>=1;  // bit 4
14           bit_in(b);  b>>=1;  // bit 5
15           bit_in(b);  b>>=1;  // bit 6
16           bit_in(b);  b>>=1;  // bit 7
17  #endif
18           return  a_;
19       }
20      [--snip--]
```

The lower block implements the straightforward idea. The program [FXT: bits/crc64-demo.cc] computes the 64-bit CRC of a single byte in both ways.

Binary words are fed in byte by byte, starting from the lower end:

```
1        uint64 word_in(uint64 w)
2        {
3            ulong k = BYTES_PER_LONG_LONG;
4            while ( k-- )  { byte_in( (uchar)w );  w>>=8; }
5            return  a_;
6        }
```

To feed in a given number of bits of a word, use the following method:

```
1        uint64 bits_in(uint64 w, uchar k)
2        // Feed in the k lowest bits of w
3        {
4            if ( k&1 )  { a_ ^= (w&1);  w >>= 1;  shift(); }
5            k >>= 1;
6
7            if ( k&1 )  { a_ ^= (w&3);  w >>= 2;  shift();  shift(); }
8            k >>= 1;
9
10           if ( k&1 )  { a_ ^= (w&15);  w >>= 4;  shift();  shift();  shift();  shift(); }
11           k >>= 1;
12
13           while ( k-- )  { byte_in( (uchar)w );  w>>=8; }
14
15           return  a_;
16       }
```

The operation is the optimized equivalent to

```
while ( k-- )  { bit_in( (uchar)w );  w>>=1; }
```

If two sequences differ in a single block of up to 64 bits, their CRCs will be different. The probability that two different sequences have the same CRC equals $2^{-64} \approx 5.42 \cdot 10^{-20}$. If that is not enough (and one does not want to write a CRC with more than 64 bits) then one can simply use two (or more) instances where different polynomials must be used. Sixteen 'random' primitive polynomials are given [FXT: bits/crc64.cc] as static class member:

```
const uint64 crc64::cc[] = {
    0x5a0127dd34af1e81ULL,  // [0]
    0x4ef12e145d0e3ccdULL,  // [1]
    0x16503f45acce9345ULL,  // [2]
    0x24e8034491298b3fULL,  // [3]
    0x9e4a8ad2261db8b1ULL,  // [4]
    0xb199aecfbb17a13fULL,  // [5]
    0x3f1fa2cc0dfbbf51ULL,  // [6]
    0xfb6e45b2f694fb1fULL,  // [7]
    0xd4597140a01d32edULL,  // [8]
    0xbd08ba1a2d621bffULL,  // [9]
    0xae2b680542730db1ULL,  // [10]
    0x8ec06ec4a8fe8f6dULL,  // [11]
    0xb89a2ecea2233001ULL,  // [12]
    0x8b996e790b615ad1ULL,  // [13]
    0x7eaef8397265e1f9ULL,  // [14]
    0xf368ae22deecc7c3ULL,  // [15]
};
```

These are taken from the list [FXT: data/rand64-hex-primpoly.txt]. Initialize multiple CRCs as follows:

```
crc64 crca( crc64::cc[0] );
crc64 crcb( crc64::cc[1] );
```

A class for 32-bit CRCs is given in [FXT: `class crc32` in bits/crc32.h]. Its usage is completely equivalent.

The CRC can easily be implemented in hardware and is, for example, used to detect errors in hard disk blocks. When a block is written its CRC is computed and stored in an extra word. When the block is read, the CRC is computed from the data and compared to the stored CRC. A mismatch indicates an error.

One property that the CRC does not have is cryptographic security. It is possible to intentionally create a data set with a prescribed CRC. With *secure hashes* (like MD5 and SHA) it is (practically) not possible to do so. Secure hashes can be used to 'sign' data. Imagine you distribute a file (for example an binary executable) over the Internet. You want to make sure that someone downloading the file (from any source) can verify that it is not an altered version (like, in the case of an executable, a malicious program). To do so you create a (secure!) hash value which you publish on your web site. Any person can verify the authenticity of the file by computing the hash and comparing it to the published version.

The cryptographic security of hash functions like MD5 and SHA is the object of ongoing research, see [263], [41], and [42].

1.32.1 Optimization via lookup tables

One can feed in an n-bit word w into the CRC in one step (instead of n steps) as follows: add w to (the CRC word) a. Save the lowest n bits of the result to a variable x. Right shift a by n bits. Add to a the entry x of an auxiliary table t. For $n = 8$ the operation can be implemented as [FXT: `class tcrc64` in bits/tcrc64.h]:

```
uint64 byte_in(uchar b)
{
    a_ ^= b;
    uint64 x =  t_[a_ & 255];
    a_ >>= 8;
    a_ ^= x;
    return  a_;
}
```

The size of the table t is $2^n = 256$ words. For $n = 1$ the table would have only two entries, zero and c, the polynomial used. Then the implementation reduces to

```
1       uint64 bit_in(uchar b)
2       {
3           a_ ^= (b&1);
4           bool s = (a_ & 1);
5           a_ >>= 1;
6           if ( 0!=s )  a_ ^= c_;  // t[0]=0; t[1]=c_;
7           return  a_;
8       }
```

which is equivalent to the `bit_in()` routine of the unoptimized CRC.

The lookup table is computed upon class initialization as follows:

```
1           for (ulong w=0; w<256; ++w)
2           {
3               set_a(0);
4               for (ulong k = 0; k<8; ++k)  bit_in( (uchar)w>>k );
5               t_[w] = a_;
6           }
```

The class can use tables of either 16 or 256 words. When a table of size 16 is used, the computation is about 6 times faster than with the non-optimized routine. A table of size 256 gives a speedup by a factor of 12. Optimization techniques based on lookup tables are often used in practical applications, both in hardware and in software, see [62].

1.32.2 Parallel CRCs

A very fast method for checksumming is to compute the CRCs for each bit of the fed-in words in parallel. An array of 64 words is used [FXT: `class pcrc64` in bits/pcrc64.h]:

```
1   template <typename Type>
2   class pcrc64
3   // Parallel computation of 64-bit CRCs for each bit of the input words.
4   // Primitive polynomial used is x^64 + x^4 + x^3 + x^2 + 1
5   {
6   public:
7       Type x_[64];    // CRC data
8       // bit(i) of x_[0], x_[1], ..., x_[63] is a 64-bit CRC
9       //   of bit(i) of all input words
10      uint pos_;      // position of constant polynomial term
11      const uint m_;  // mask to compute mod 64
12
```

Upon initialization all words are set to all ones:

```
1   public:
2       pcrc64()
3           : m_(63)
4       {
5           reset();
6       }
7
8       ~pcrc64()  { ; }
9
10      void reset()
11      {
12          pos_ = 0;
13          Type ff = 0;  ff = ~ff;
14          for (uint k=0; k<64; ++k)  x_[k] = ff;
15      }
```

The cyclic shift of the array is avoided by working modulo 64 when feeding in words:

```
1       void word_in(Type w)
2       {
3           uint p = pos_;
4           pos_ = (p+1) & m_;
5           uint h = (p-1) & m_;
6           Type a = x_[p & m_];  // 0
7           p += 2;
```

```
 8          a ^= x_[p & m_];  // 2
 9          ++p;
10          a ^= x_[p & m_];  // 3
11          ++p;
12          a ^= x_[p & m_];  // 4
13          x_[h] = a ^ w;
14      }
```

The algorithm corresponds to the Fibonacci setup of the linear feedback shift registers (see section 39.2 on page 860). There is no primitive trinomial with degree a multiple of eight so we use the pentanomial $x^{64} + x^4 + x^3 + x^2 + 1$. With an array size where a primitive trinomial exists the modulo computations would be more expensive. A unrolled routine can be used to feed in multiple words:

```
1       void words_in(Type *w, ulong n)
2       {
3           if ( n&1 )  { word_in(w[0]); ++w; }
4           n >>= 1;
5
6           if ( n&1 )  { word_in(w[0]); word_in(w[1]); w+=2; }
7           n >>= 1;
8
9           for (ulong k=0; k<n; ++k)
10          {
11              word_in(w[0]);
12              word_in(w[1]);
13              word_in(w[2]);
14              word_in(w[3]);
15              w += 4;
16          }
17      }
```

The program [FXT: bits/pcrc64-demo.cc] feeds the numbers up to a given value into a `pcrc64<uint>`:

```
int main()
{
    Type n = 32768;
    pcrc64<Type> P;
    for (Type k=0; k<n; ++k)  P.word_in(k);

    // print array P.x_[] here
}
```

This rather untypical type of input data illustrates the independence of the bits in the array `x_[]`:

```
     0:  ................1..1.1.11.111.
     1:  ................i.ii.i.1ii1.1ii.
     2:  ................i...ii...1...i.
     3:  ................i..iii....ii.
     4:  ................i..i.1ii1.ii..
     5:  ................ii...i.1...i.i
     6:  ................i.i.1111.i.ii.
     7:  ................i.i.11.i...i..
     8:  ................i...i.iiii1.
     9:  ................iiii.i.i.i..
    10:  ..................iii...i..i..i
    11:  ................iiii.i...i.i.i
    12:  ................ii..i.ii.ii..i
    13:  ................ii.i.i..ii.iiii
   [--snip--]
    59:  ................1...1.11.11.1.1
    60:  ................i.i..ii.1ii..i.i
    61:  ................1..iiii...ii.
    62:  ................i..i..ii...i..
    63:  ................iii..i.iiii.iii
```

The implementation can process about 2 GB of data per second when 64-bit types are used, 1 GB/s with 32-bit types, 500 MB/sec with 16-bit types, and about 230 MB/sec with 8-bit types.

1.33 Permutations via primitives

We give two methods to specify permutations of the bits of a binary word via one or more control words. The methods are suggestions for machine instructions that can serve as primitives for permutations of the bits of a word.

A restricted method

Generalizing the symbolic-powering idea given in section 1.19 on page 49 we can specify a subset of all permutations by selecting bit-blocks of the masks (for 32-bit words)

```
................1111111111111111
........11111111........11111111
....1111....1111....1111....1111
..11..11..11..11..11..11..11..11
.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1
```

Subsets of the blocks of the masks can be determined with the bits of a 32-bit word by considering the highest bit of each block:

```
................1................    bits  15 ...
........1.......1.......1.......    bits   7 ...
....1.....1.....1.....1.....1...    bits   3  11 ...
..1...1...1...1...1...1...1...1.    bits   1   5   9  13 ..
.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1    bits   0   2   4   6   8  10  12  14 ...
```

Thereby one can use all (except for the highest) bits of a word to select the blocks where the bits defined by the block and those left to it should be swapped. A implementation of the implied algorithm can be found in [FXT: bits/bitperm1-demo.cc]. Arrays are used to give more readable code:

```
1   void perm1(uchar *a, ulong ldn, const uchar *x)
2   // Permute a[] according to the 'control word' x[].
3   // The length of a[] must be 2**ldn.
4   {
5       long n = 1L<<ldn;
6       for (long s=n/2; s>0; s/=2)
7       {
8           for (long k=0; k<n; k+=s+s)
9           {
10              if ( x[k+s-1]!='0' )
11              {
12                  // swap regions [a+k,...,a+k+s-1], [a+k+s,...,a+k+2*s-1]:
13                  swap(a+k, a+k+s, s);
14              }
15          }
16      }
17  }
```

No attempt has been made to optimize or parallelize the algorithm. We just explore how useful a machine instruction for the permutation of bits would be.

The program uses a fixed size of 16 bits, an 'x' is printed whenever the corresponding bit is set:

```
a=0123456789ABCDEF   bits of the input word
x=0010011000110110   control word
   8:    7
   4:    3  11x
   2:    1   5x  9  13x
   1:    0   2x  4   6x  8  10x 12  14x
a=01326754CDFEAB98   result
```

The control word used leads to the Gray code permutation (see 2.8 on page 105). Assume we use words with N bits. We cannot (for $N > 2$) obtain all $N!$ permutations as we can choose between only 2^{N-1} control words. Now set the word length to $N := 2^n$. The reachable permutations are those where the intervals $[k \cdot 2^j, \ldots, (k+1) \cdot 2^j - 1]$ contain all numbers $[p \cdot 2^j, \ldots, (p+1) \cdot 2^j - 1]$ for all $j \le n$ and $0 \le k < 2^{n-j}$, choosing p for each interval arbitrarily ($0 \le p < 2^{n-j}$). For example, the lower half of the permuted array must contain a permutation of either the lower or the upper half ($j = n-1$) and each pair a_{2y}, a_{2y+1} must contain two elements $2z$, $2z+1$ ($j = 1$). The bit-reversal can be obtained using a control word with all bits set. Alas, the (important!) zip permutation (bit-zip, see section 1.15 on page 39) is unreachable.

The inverse permutation is achieved by changing a single line:

```
for (long s=1; s<n; s+=s)
```

A machine instruction could choose between the two routines via the highest bit in the control word.

A general method

All permutations of $N = 2^n$ elements can be obtained using n control words of N bits. Assume we have a machine instruction that collects bits according to a control word. An eight bit example:

```
a = abcdefgh   input data
x = ..1.11.1   control word (dots for zeros)
      cefh     bits of a where x has a one
    abdg       bits of a where x has a zero
    abdgcefh   result, bits separated according to x
```

We need n such instructions that work on all length-2^k sub-words for $1 \le k \le n$. For example, the instruction working on half words of a 16-bit word would work as

```
a = abcdefgh ABCDEFGH   input data
x = ..1.11.1 1111....   control word (dots for zeros)
      cefh      ABCD    bits of a where x has a one
    abdg     EFGH       bits of a where x has a zero
    abdgcefh EFGHABCD   result, bits separated according to x
```

Note the bits of the different sub-words are not mixed. Now all permutations can be reached if the control word for the 2^k-bit sub-words have exactly 2^{k-1} bits set in all ranges $[j \cdot 2^k, \ldots, (j+1) \cdot 2^k]$.

A control word together with the specification of the instruction used defines the action taken. The following leads to a swap of adjacent bit pairs

```
1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.    k= 1  (2-bit sub-words)
```

while this

```
1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1.    k= 5 (32 bit sub-words)
```

results in gathering the even and odd indexed bits in the halfwords.

A complete set of permutation primitives for 16-bit words and their effect on a symbolic array of bits (split into groups of four elements for readability) is

```
                                0123 4567 89ab cdef
11111111........    k= 4  ==> 89ab cdef 0123 4567
1111....1111....    k= 3  ==> cdef 89ab 4567 0123
11..11..11..11..    k= 2  ==> efcd ab89 6745 2301
1.1.1.1.1.1.1.1.    k= 1  ==> fedc ba98 7654 3210
```

The top primitive leads to a swap of the left and right half of the bits, the next to a swap of the halves of the half words and so on. The obtained permutation is array reversal. Note that we use array notation (least index left) here.

The resulting permutation depends on the order in which the primitives are used. Starting with full words

```
                                    0123 4567 89ab cdef
1.1. 1.1. 1.1. 1.1.  k= 4  ==>    1357 9bdf 0246 8ace
1.1. 1.1. 1.1. 1.1.  k= 3  ==>    37bf 159d 26ae 048c
1.1. 1.1. 1.1. 1.1.  k= 2  ==>    7f3b 5d19 6e2a 4c08
1.1. 1.1. 1.1. 1.1.  k= 1  ==>    f7b3 d591 e6a2 c480
```

The result is different when starting with 2-bit sub-words:

```
                                    0123 4567 89ab cdef
1.1. 1.1. 1.1. 1.1.  k= 1  ==>    1032 5476 98ba dcfe
1.1. 1.1. 1.1. 1.1.  k= 2  ==>    0213 4657 8a9b cedf
1.1. 1.1. 1.1. 1.1.  k= 3  ==>    2367 0145 abef 89cd
1.1. 1.1. 1.1. 1.1.  k= 4  ==>    3715 bf9d 2604 ae8c
```

There are $\binom{2z}{z}$ possibilities to have z bits set in a $2z$-bit word. There are 2^{n-k} length-2^k sub-words in a 2^n-bit word so the number of valid control words for that step is

$$\left[\binom{2^k}{2^{k-1}}\right]^{2^{n-k}}$$

The product of the number of valid words in all steps gives the number of permutations:

$$(2^n)! \;=\; \prod_{k=1}^{n} \left[\binom{2^k}{2^{k-1}}\right]^{2^{n-k}} \tag{1.33-1}$$

1.34 CPU instructions often missed

Essential

- Bit-shift and bit-rotate instructions that work properly for shifts greater or equal to the word length: the shift instruction should zero the word, the rotate instruction should take the shift modulo word length. The C-language standards leave the results for these operations undefined and compilers simply emit the corresponding assembler instructions. The resulting CPU dependent behavior is both a source of errors and makes certain optimizations impossible.
- A bit-reverse instruction. A fast byte-swap mitigates the problem, see section 1.14 on page 34.
- Instructions that return the index of highest or lowest set bit in a word.
- Fast conversion from integer to float and double (both directions).
- A fused multiply-add instruction for floats.
- Instructions for the multiplication of complex floating point numbers. For example, with 128-bit (SIMD) registers and 64-bit floats:

```
// Here: R1 == A, B;  R2 == C, D;
CMUL R1, R2;
// Here: R2 == A*C-B*D, A*D+B*C;
```

- A sum-diff instruction, such as:

```
// Here:  R1 == A;  R2 == B;
SUMDIFF R1, R2;
// Here:  R1 == A+B;  R2 == A-B;
```

  It serves as a primitive for fast orthogonal transforms.

- An instruction to swap registers. Even better, a conditional version of that.

Nice to have

- A parity bit for the complete machine word. The parity of a word is the number of bits modulo two, not the complement of it. Even better would be an instruction for the inverse Gray code, see section 1.16 on page 40.
- A bit-count instruction, see section 1.8 on page 19. This would also give the parity at bit zero.
- A random number generator: LHCAs (see section 39.7 on page 866) may be candidates.
- A conditional version of more than just the move instruction, possibly as an instruction prefix.
- An instruction to detect zero bytes in a word, see section 1.22 on page 58. The C-convention is to use a zero byte as string terminator. Performance of the string related functions in the C-library could thereby be increased significantly. Ideally the instruction should exist for different word sizes: 4-byte, 8-byte and 16-byte (possibly using SIMD extensions).
- A bit-zip and a bit-unzip instruction, see section 1.15 on page 39. Note this is polynomial squaring over GF(2).
- A bit-gather and a bit-scatter instruction, see [FXT: bits/bitgather.h] and [FXT: bits/bitseparate.h]. This would include bit-zip and its inverse.
- Primitives for permutations of bits, see section 1.33 on page 88.
- Multiplication corresponding to XOR as addition. That is, a multiplication without carries, the one used for polynomials over GF(2). See section 38.1 on page 815 and [FXT: `bitpol_mult()` in bpol/bitpol-arith.h].

Chapter 2

Permutations

In this chapter we first study several special permutations that are used as part of other algorithms like the revbin permutation used in radix-2 FFT algorithms. Then permutations in general together with the operations on them, like composition and inversion, are described.

Algorithms for the generation of all permutations of a given number of objects are given in chapter 10. The connection to mixed radix numbers in factorial base is given in section 10.3.

2.1 The revbin permutation

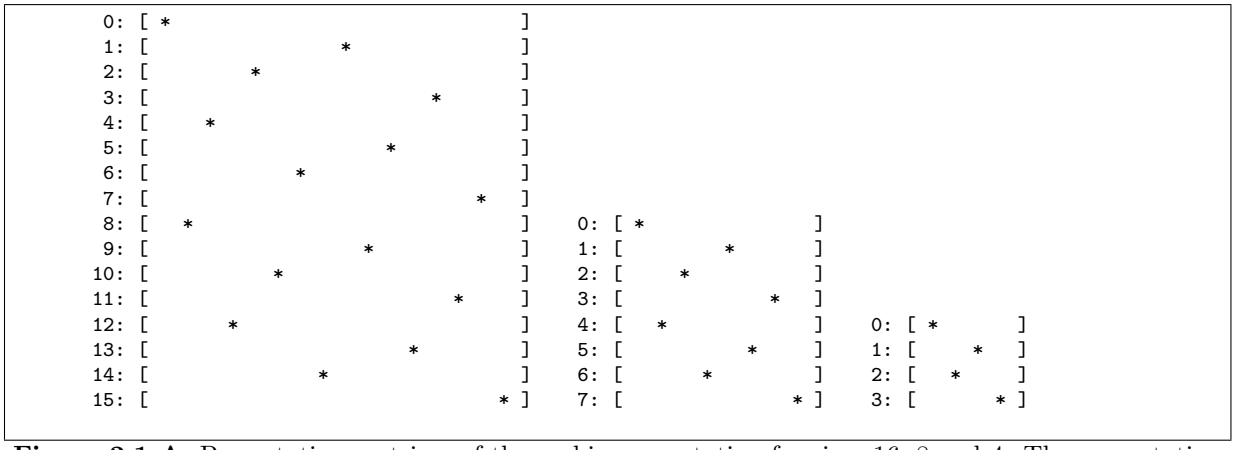

Figure 2.1-A: Permutation matrices of the revbin permutation for sizes 16, 8 and 4. The permutation is self-inverse.

The permutation that swaps elements whose binary indices are mutual reversals is called *revbin permutation* (sometimes also *bit-reversal-* or *bitrev* permutation). For example, for length $n = 256$ the element with index $x = 43_{10} = 00101011_2$ is swapped with the element whose index is $\tilde{x} = 11010100_2 = 212_{10}$. Note that $\tilde{x}$ depends on both x and on n. Pseudo code for a naive implementation is

```
1   procedure revbin_permute(a[], n)
2   // a[0..n-1] input,result
3   {
4       for x:=0 to n-1
5       {
6           r := revbin(x, n)
7           if r>x then  swap(a[x], a[r])
8       }
9   }
```

The condition `r>x` before the `swap()` statement makes sure that the swapping is not undone later when the loop variable `x` has the value of the present `r`. The key ingredient for a fast permutation routine is the observation that we only need to update as bit-reversed values. Given $\tilde{x}$ we can compute $\widetilde{x+1}$ efficiently. Methods to do this are given in section 1.14.3. A faster routine will be of the form

```
1   procedure revbin_permute(a[], n)
2   // a[0..n-1] input,result
3   {
4       if n<=2  return
5       r := 0  // the reversed 0
6       for x:=1 to n-1
7       {
8           r := revbin_upd(r, n/2)  // inline me
9           if r>x then  swap(a[x], a[r])
10      }
11  }
```

A method to generate all revbin pairs via linear feedback shift registers is given in section 39.3 on page 861.

About $(n-\sqrt{n})/2$ `swap()`-statements will be executed with the revbin permutation of n elements. That is, almost every element is moved for large n, as there are only few numbers with symmetric bit patterns:

n:	2 # swaps	# symm. pairs
2:	0	2
4:	2	2
8:	4	4
16:	12	4
32:	24	8
64:	56	8
2^{10}:	992	32
2^{20}:	$0.999 \cdot 2^{20}$	2^{10}
∞ :	$n-\sqrt{n}$	$\sqrt{n}$

The sequence

```
0, 2, 4, 12, 24, 56, 112, 238, 480, 992, 1980, 4032, 8064, 16242, 32512, 65280, ...
```

is entry A045687 of [240].

Optimizations using the symmetries of the permutation

The following symmetry can be used for further optimization: if for even $x < \frac{n}{2}$ there is a swap (for the pair x, $\tilde{x}$) then there is also a swap for the pair $n-1-x$, $n-1-\tilde{x}$. As $x < \frac{n}{2}$ and $\tilde{x} < \frac{n}{2}$ one has $n-1-x > \frac{n}{2}$ and $n-1-\tilde{x} > \frac{n}{2}$. That is, the swaps are independent. A routine that uses these observations is

```
1   procedure revbin_permute(a[], n)
2   {
3       if n<=2  return
4       nh := n/2
5       r := 0  // the reversed 0
6       x := 1
7       while x<nh
8       {
9           // x odd:
10          r := r + nh
11          swap(a[x], a[r])
12          x := x + 1
13
14          // x even:
15          r := revbin_upd(r,n/2)  // inline me
16          if r>x then
17          {
18              swap(a[x], a[r])
19              swap(a[n-1-x], a[n-1-r])
20          }
21          x := x + 1
22      }
23  }
```

The code above can be used to derive an optimized version for zero padded data:

```
procedure revbin_permute0(a[], n)
{
    if n<=2  return
    nh := n/2
    r := 0  // the reversed 0
    x := 1
    while x<nh
    {
        // x odd:
        r := r + nh
        a[r] := a[x]
        a[x] := 0
        x := x + 1

        // x even:
        r := revbin_upd(r, n)  // inline me
        if r>x then  swap(a[x], a[r])
        // Omit swap of a[n-1-x] and a[n-1-r] as these are zero
        x := x + 1
    }
}
```

One can carry the scheme further, distinguishing whether $x \bmod 4 = 0$, 1, 2 or 3. This is done in the actual C++ implementation [FXT: `revbin_permute()` in perm/revbinpermute.h]. The parameters

```
#define  RBP_SYMM  4  // 1, 2, 4 (default is 4)
#define  FAST_BIT_SCAN  // define if machine has fast bit scan
```

determine how much of the symmetry is used (`RBP_SYMM`) and which flavor of the revbin-update routine is chosen (`FAST_BIT_SCAN`). With a fast bit-scan instruction the table driven version is slightly faster. We further define a convenient macro to swap elements:

```
#define  idx_swap(k, r)  { ulong kx=(k), rx=(r);  swap2(f[kx], f[rx]); }
```

The main routine uses unrolled versions of the revbin permutation for small values of n. These are given in [FXT: perm/shortrevbinpermute.h]. For example, the unrolled routine for $n = 16$ is

```
template <typename Type>
inline void revbin_permute_16(Type *f)
{
    swap2(f[1], f[8]);
    swap2(f[2], f[4]);
    swap2(f[3], f[12]);
    swap2(f[5], f[10]);
    swap2(f[7], f[14]);
    swap2(f[11], f[13]);
}
```

The code was generated with the program [FXT: perm/cycles-demo.cc], see section 2.11.3 on page 114. The routine `revbin_permute_le_64(f,n)` that is called for $n \leq 64$ selects the correct routine for the parameter n:

```
template <typename Type>
void revbin_permute(Type *f, ulong n)
{
    if ( n<=64 )
    {
        revbin_permute_le_64(f, n);
        return;
    }
}
```

If the table driven update is used, the table has to be initialized and, depending on the amount of symmetry used, also some auxiliary variables:

```
    const ulong nh = (n>>1);
#ifdef FAST_BIT_SCAN
    make_revbin_upd_tab(nh);
#endif

#if  ( RBP_SYMM >= 2 )
    const ulong n1  = n - 1;    // = 11111111
#if  ( RBP_SYMM >= 4 )
    const ulong nx1 = nh - 2;   // = 01111110
    const ulong nx2 = n1 - nx1; // = 10111101
#endif //  ( RBP_SYMM >= 4 )
```

```
#endif //  ( RBP_SYMM >= 2 )
```

In what follows we set RBP_SYMM to 4 and omit the corresponding preprocessor statements. The cryptic and powerful main loop is

```
    ulong k = 0,  r = 0;
    while ( k<n/RBP_SYMM  )  // n>=16, n/2>=8, n/4>=4
    {
        // ----- k%4 == 0:
        if ( r>k )
        {
            idx_swap(k, r);  // <nh, <nh 11
            idx_swap(n1^k, n1^r);  // >nh, >nh 00
            idx_swap(nx1^k, nx1^r);  // <nh, <nh 11
            idx_swap(nx2^k, nx2^r);  // >nh, >nh 00
        }

        r ^= nh;
        ++k;

        // ----- k%4 == 1:
        if ( r>k )
        {
            idx_swap(k, r);  // <nh, >nh 10
            idx_swap(n1^k, n1^r);  // >nh, <nh 01
        }

#ifdef FAST_BIT_SCAN
        r = revbin_tupd(r, k);
#else
        r = revbin_upd(r, nh);
#endif
        ++k;

        // ----- k%4 == 2:
        if ( r>k )
        {
            idx_swap(k, r);  // <nh, <nh 11
            idx_swap(n1^k, n1^r); // >nh, >nh 00
        }

        r ^= nh;
        ++k;

        // ----- k%4 == 3:
        if ( r>k )
        {
            idx_swap(k, r);    // <nh, >nh 10
            idx_swap(nx1^k, nx1^r);   // <nh, >nh 10
        }

#ifdef FAST_BIT_SCAN
        r = revbin_tupd(r, k);
#else
        r = revbin_upd(r, nh);
#endif
        ++k;
    }
}  // end revbin_permute()
```

It turns out that the routine takes, for large n, about six times of the simple `reverse()` operation that swaps elements k with $n-k-k$. Much of the time is spend waiting for memory which suggests that further optimizations would best be attempted with special (machine) instructions to bypass the cache or with non-temporal writes.

The routine [FXT: `revbin_permute0()` in perm/revbinpermute0.h] is a specialized version optimized for zero padded data. Some memory access can be avoided for that case. For example, revbin-pairs with both indices larger than $n/2$ need no processing at all. Therefore the routine is faster than the general version.

If, for complex data, one works with separate arrays for the real and imaginary parts one can remove half of the bookkeeping as follows:

```
procedure revbin_permute(a[], b[], n)
{
    if n<=2  return
    r := 0  // the reversed 0
```

```
5       for x:=1 to n-1
6       {
7           r := revbin_upd(r, n/2)  // inline me
8           if r>x then
9           {
10              swap(a[x], a[r])
11              swap(b[x], b[r])
12          }
13      }
14  }
```

If both real and imaginary part fit into level-1 cache the method can lead to a speedup. However, for large arrays the routine can be drastically slower than two separate calls of the simple method: typically the real and imaginary element for the same index lie apart in memory by a power of two, leading to one hundred percent cache miss for large array sizes.

2.2 The radix permutation

The *radix permutation* is the generalization of the revbin permutation to arbitrary radices. Pairs of elements are swapped when their indices, written in radix r are reversed. For example, in radix 10 and $n = 1000$ the elements with indices 123 and 321 will be swapped. The radix permutation is self-inverse.

C++ code for the radix r permutation of the array `f[]` is given in [FXT: perm/radixpermute.h]. The routine must be called with n a perfect power of the radix r. Using radix $r = 2$ gives the revbin permutation.

```
1   extern ulong radix_permute_nt[];  // == 9, 90, 900, ...  for r=10
2   extern ulong radix_permute_kt[];  // == 1, 10, 100, ...  for r=10
3   #define  NT  radix_permute_nt
4   #define  KT  radix_permute_kt
5
6   template <typename Type>
7   void radix_permute(Type *f, ulong n, ulong r)
8   {
9       ulong x = 0;
10      NT[0] = r-1;
11      KT[0] = 1;
12      while ( 1 )
13      {
14          ulong z = KT[x] * r;
15          if ( z>n )  break;
16          ++x;
17          KT[x] = z;
18
19          NT[x] = NT[x-1] * r;
20      }
21      // here: n == p**x
22
23      for (ulong i=0, j=0;  i < n-1;  i++)
24      {
25          if ( i<j )  swap2(f[i], f[j]);
26
27          ulong t = x - 1;
28          ulong k = NT[t];  // =^=  k = (r-1) * n / r;
29
30          while ( k<=j )
31          {
32              j -= k;
33              k = NT[--t];  // =^=  k /= r;
34          }
35
36          j += KT[t]; // =^=  j += (k/(r-1));
37      }
38  }
```

2.3 In-place matrix transposition

Transposing a matrix is easy when it is not done in-place. The simple routine [FXT: `transpose()` in aux2/transpose.h] does the job:

```
template <typename Type>
void transpose(const Type * restrict f, Type * restrict g, ulong nr, ulong nc)
// Transpose nr x nc matrix f[]  into an nc x nr matrix g[].
{
    for (ulong r=0; r<nr; r++)
    {
        ulong isrc = r * nc;
        ulong idst = r;
        for (ulong c=0; c<nc; c++)
        {
            g[idst] = f[isrc];
            isrc += 1;
            idst += nr;
        }
    }
}
```

Matters get more complicated for the in-place equivalent. We have to find the cycles (see section 2.11 on page 112) of the underlying permutation. To transpose a $n_r \times n_c$- matrix first identify the position i of the entry in row r and column c:

$$i = r \cdot n_c + c \tag{2.3-1}$$

After the transposition the element will be at position i' in the transposed $n'_r \times n'_c$- matrix

$$i' = r' \cdot n'_c + c' \tag{2.3-2}$$

Obviously, $r' = c$, $c' = r$, $n'_r = n_c$ and $n'_c = n_r$, so:

$$i' = c \cdot n_r + r \tag{2.3-3}$$

Multiply the last equation by n_c

$$i' \cdot n_c = c \cdot n_r \cdot n_c + r \cdot n_c \tag{2.3-4}$$

With $n := n_r \cdot n_c$ and $r \cdot n_c = i - c$ we get

$$i' \cdot n_c = c \cdot n + i - c \tag{2.3-5}$$

$$i = i' \cdot n_c - c \cdot (n-1) \tag{2.3-6}$$

Take the equation modulo $n - 1$ to get

$$i \equiv i' \cdot n_c \mod (n-1) \tag{2.3-7}$$

That is, the transposition moves the element $i = i' \cdot n_c$ to position i'. Multiply by n_r to get the inverse:

$$i \cdot n_r \equiv i' \cdot n_c \cdot n_r \tag{2.3-8}$$

$$i \cdot n_r \equiv i' \cdot (n - 1 + 1) \tag{2.3-9}$$

$$i \cdot n_r \equiv i' \tag{2.3-10}$$

That is, element i will be moved to $i' = i \cdot n_r \mod (n-1)$.

The routine [FXT: `transpose()` in aux2/transpose.h] uses the a bit-array to keep track of the elements that have been processed so far:

```
#define  SRC(k)  (((unsigned long long)(k)*nc)%n1)

template <typename Type>
void transpose(Type *f, ulong nr, ulong nc, bitarray *ba=0)
```

```
5   // In-place transposition of an  nr X nc  array
6   //   that lies in contiguous memory.
7   {
8       if ( 1>=nr )  return;
9       if ( 1>=nc )  return;
10
11      if ( nr==nc )  transpose_square(f, nr);
12      else
13      {
14          const ulong n1 = nr * nc - 1;
15
16          bitarray *tba = 0;
17          if ( 0==ba )  tba = new bitarray(n1);
18          else          tba = ba;
19          tba->clear_all();
20
21          for (ulong k=1;  k<n1;  k=tba->next_clear(++k) )  //  0 and n1 are fixed points
22          {
23              // do a cycle:
24              ulong ks = SRC(k);
25              ulong kd = k;
26              tba->set(kd);
27              Type t = f[kd];
28              while ( ks != k )
29              {
30                  f[kd] = f[ks];
31                  kd = ks;
32                  tba->set(kd);
33                  ks = SRC(ks);
34              }
35              f[kd] = t;
36          }
37
38          if ( 0==ba )  delete tba;
39      }
40  }
```

Note that one should take care of possible overflows in the calculation $i \cdot n_c$. For the case that n is a power of two (and so are both n_r and n_c) the multiplications modulo $n-1$ are cyclic shifts. Thus any overflow can be avoided and the computation is also significantly cheaper. A C++ implementation is given in [FXT: aux2/transpose2.h].

2.4 Revbin permutation and matrix transposition *

```
Rotate left by 3 positions:
  [ 1 2 3 4 5 6 7 8 ]  original array
  [ 3 2 1 4 5 6 7 8 ]  reverse first 3 elements
  [ 3 2 1 8 7 6 5 4 ]  reverse last 8-3=5 elements
  [ 4 5 6 7 8 1 2 3 ]  reverse whole array


Rotate right by 3 positions:
  [ 1 2 3 4 5 6 7 8 ]  original array
  [ 5 4 3 2 1 6 7 8 ]  reverse first 8-3=5 elements
  [ 5 4 3 2 1 8 7 6 ]  reverse last 3 elements
  [ 6 7 8 1 2 3 4 5 ]  reverse whole array
```

Figure 2.4-A: Rotation of an length-8 array by 3 positions to the left (top), and right (bottom).

We give a method to rotate an (length-n) array by s positions (left or right) that does not require any scratch space. The trick is to use `reverse()` three times as in the following [FXT: `rotate_left()` in perm/rotate.h]:

```
1   template <typename Type>
2   void rotate_left(Type *f, ulong n, ulong s)
3   // Rotate towards element #0
4   // Shift is taken modulo n
5   {
```

```
6       if ( s==0 )  return;
7       if ( s>=n )
8       {
9           if (n<2) return;
10          s %= n;
11      }
12
13      reverse(f,   s);
14      reverse(f+s, n-s);
15      reverse(f,   n);
16  }
```

The technique is usually called the *triple reversion trick*. For example left-rotating an 8-element array by 3 positions is achieved by the steps shown in figure 2.4-A (top). Similarly for the other direction. A right rotation of an n-element array by s positions is identical to a left rotation by $n - s$ positions (bottom of figure 2.4-A).

```
1   template <typename Type>
2   void rotate_right(Type *f, ulong n, ulong s)
3   // Rotate away from element #0
4   // Shift is taken modulo n
5   {
6       if ( s==0 )  return;
7       if ( s>=n )
8       {
9           if (n<2) return;
10          s %= n;
11      }
12
13      reverse(f,     n-s);
14      reverse(f+n-s, s);
15      reverse(f,     n);
16  }
```

When transposing an $n_r \times n_c$ matrix whose size is a power of two (thereby both n_r and n_c are also powers of two) the above mentioned rotation is done with the *indices* (written in base two) of the elements. We know how to do a permutation that reverses the complete indices and reversing a few bits at the least significant end is not any harder:

```
1   template <typename Type>
2   void revbin_permute_rows(Type *f, ulong ldn, ulong ldnc)
3   // Revbin-permute the length 2**ldnc rows  of f[0..2**ldn-1]
4   // (f[] considered as an 2**(ldn-ldnc) x 2**ldnc matrix)
5   {
6       ulong n = 1UL<<ldn;
7       ulong nc = 1UL<<ldnc;
8       for (ulong k=0; k<n; k+=nc)  revbin_permute(f+k, nc);
9   }
```

The routine to transpose a matrix is

```
1   template <typename Type>
2   void transpose_by_rbp(Type *f, ulong ldn, ulong ldnc)
3   // Transpose f[] considered as an 2**(ldn-ldnc) x 2**ldnc matrix
4   {
5       revbin_permute_rows(f, ldn, ldnc);
6       ulong n = 1UL<<ldn;
7       revbin_permute(f, n);
8       revbin_permute_rows(f, ldn, ldn-ldnc);  // ... that is, columns
9   }
```

```
                v v v v v     v v v v         <--= want to swap these blocks
  [ 0 1 2 3 4   a b c d e 7 8 w x y z  N N ]  original array
  [ 0 1 2 3 4   e d c b a 7 8 w x y z  N N ]  reverse first block
  [ 0 1 2 3 4   e d c b a 8 7 w x y z  N N ]  reverse range between blocks
  [ 0 1 2 3 4   e d c b a 8 7 z y x w  N N ]  reverse second block
  [ 0 1 2 3 4   w x y z 7 8 a b c d e  N N ]  reverse whole range
                ^ ^ ^ ^     ^ ^ ^ ^ ^         <--= the swapped blocks
```

Figure 2.4-B: Swapping the blocks `[a b c d e]` and `[w x y z]` via 4 reversals.

The triple-reversion trick can also be used to swap two blocks in an array: first reverse the three ranges (first blocks, range between block, last block), then reverse the range that consists of all three. This is the *quadruple reversion trick*. The corresponding code is given in [FXT: perm/swapblocks.h]:

```
template <typename Type>
void swap_blocks(Type *f, ulong x1, ulong n1, ulong x2, ulong n2)
// Swap the blocks starting at indices x1 and x2
// n1 and n2 are the block lengths
{
    if ( x1>x2 ) { swap2(x1,x2); swap2(n1,n2); }
    f += x1;
    x2 -= x1;
    ulong n = x2 + n2;
    reverse(f, n1);
    reverse(f+n1, n-n1-n2);
    reverse(f+x2, n2);
    reverse(f, n);
}
```

The elements before `x1` and after `x2+n2` are not accessed. An example [FXT: perm/swap-blocks-demo.cc] is shown in figure 2.4-B. The effect of `swap_blocks(f, x1, n1, x2, n2)` can be undone with the statement

```
swap_blocks(f, x1, n2, x2+n2-n1, n1)
```

2.5 The zip permutation

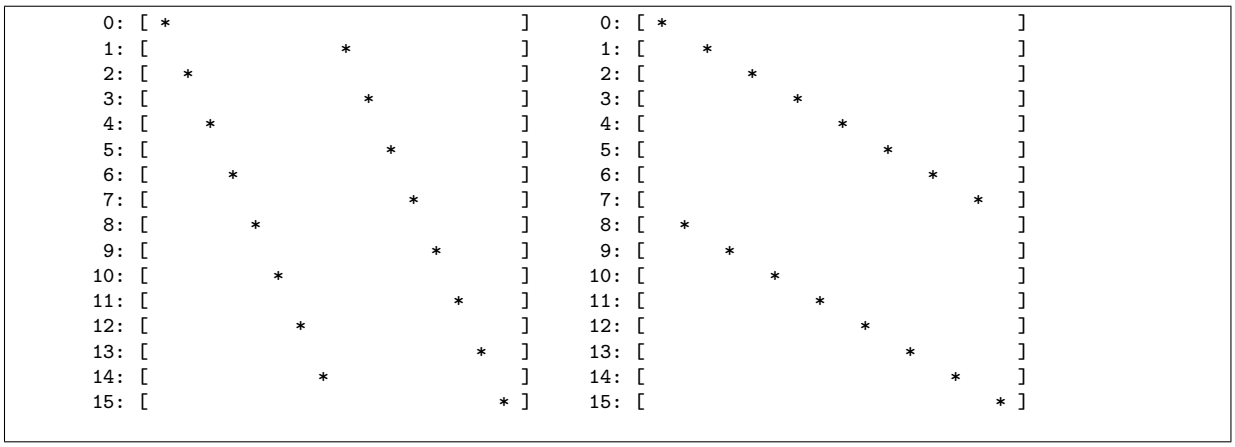

Figure 2.5-A: Permutation matrices of the zip permutation (left) and its inverse, the unzip permutation (right). The zip permutation moves the lower half of the array to the even indices and the upper half to the odd indices.

The *zip permutation* moves the elements from the lower half to the even indices and the elements from the upper half to the odd indices. Symbolically,

```
[ a b c d A B C D ]  |--> [ a A b B c C d D ]
```

The size of the array must be even. A routine for the permutation is

```
template <typename Type>
void zip(const Type * restrict f, Type * restrict g, ulong n)
{
    ulong nh = n/2;
    for (ulong k=0,  k2=0;  k<nh;  ++k, k2+=2)  g[k2] = f[k];
    for (ulong k=nh, k2=1;  k<n;   ++k, k2+=2)  g[k2] = f[k];
}
```

When the array size is a power of two we can use a special case of the 'transposition by revbin permutation' idea to do the operation in-place [FXT: `zip()` in perm/zip.h]:

```
template <typename Type>
```

```
2   void zip(Type *f, ulong n)
3   {
4       ulong nh = n/2;
5       revbin_permute(f, nh);  revbin_permute(f+nh, nh);
6       revbin_permute(f, n);
7   }
```

If we have a type `Complex` consisting of two `double`s lying contiguous in memory we can optimize the procedure:

```
1   void zip(double *f, long n)
2   {
3       revbin_permute(f, n);
4       revbin_permute((Complex *)f, n/2);
5   }
```

The inverse of the zip permutation is the *unzip permutation*. We only give the in-place version based on the revbin permutation, the array size must be a power of two:

```
1   template <typename Type>
2   void unzip(Type *f, ulong n)
3   {
4       ulong nh = n/2;
5       revbin_permute(f, n);
6       revbin_permute(f, nh);  revbin_permute(f+nh, nh);
7   }
```

The routine can for the type `double` again be optimized as

```
1   void unzip(double *f, long n)
2   {
3       revbin_permute((Complex *)f, n/2);
4       revbin_permute(f, n);
5   }
```

Connection to matrix transposition

For arrays whose size n is not a power of two the in-place zip permutation can be obtained by transposing the data as a $2 \times n/2$ matrix:

```
transpose(f, 2, n/2);  // =^= zip(f, n)
```

The routines for in-place transposition are given in section 2.3 on page 98. The inverse is clearly obtained by transposing the data as a $n/2 \times 2$ matrix:

```
transpose(f, n/2, 2);  // =^= unzip(f, n)
```

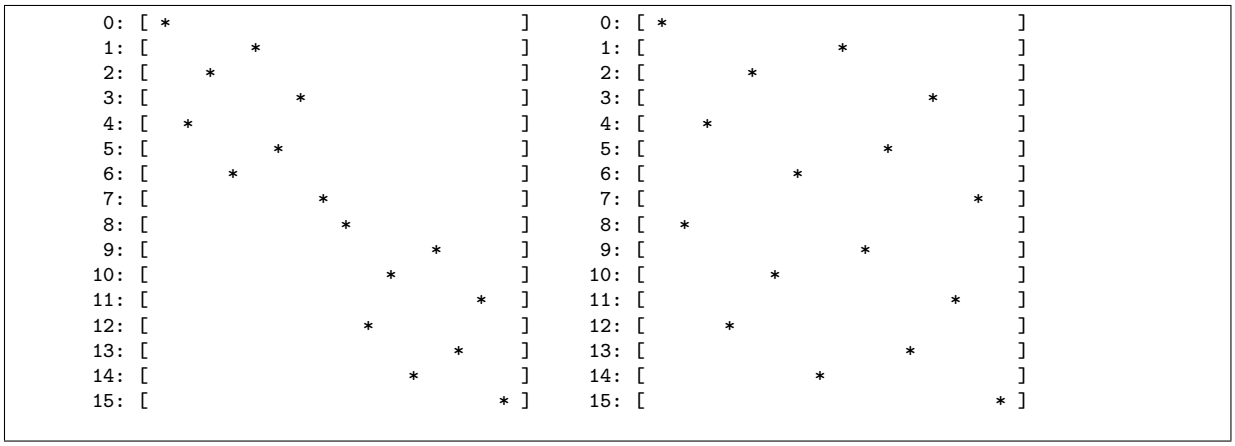

Figure 2.5-B: Revbin permutation matrices that, when multiplied together, give the zip permutation and its inverse. Let L and R be the permutations given on the left and right side, respectively. Then $Z = R\,L$ and $Z^{-1} = L\,R$.

While the above mentioned technique is usually *not* a gain for doing a transposition it may be used to speed up the revbin permutation itself. We operator-ize the idea to see how. Let R be the revbin permutation `revbin_permute()`, $T(n_r, n_c)$ the transposition of the $n_r \times n_c$ matrix and $R(n_c)$ the operation done by `revbin_permute_rows()` (see section 2.4 on page 99). Then

$$T(n_r, n_c) \quad = \quad R(n_r) \cdot R \cdot R(n_c) \tag{2.5-1}$$

The R-operators are their own inverses while T is in general not self inverse[1].

$$R \quad = \quad R(n_r) \cdot T(n_r, n_c) \cdot R(n_c) \tag{2.5-2}$$

There is a degree of freedom in this formula: for fixed $n = n_r \times n_c$ one can choose one of n_r and n_c (only their product is given).

2.6 The reversed zip permutation

A permutation closely related to the zip permutation is the *reversed zip permutation*. It moves the lower half of an array to the even indices and the upper half to the odd indices in reversed order. Symbolically,

```
[ a b c d A B C D ]  |--> [ a D b C c B d A ]
```

A C++ routine is [FXT: `zip_rev()` in perm/ziprev.h]:

```
1  template <typename Type>
2  void zip_rev(const Type * restrict x, Type * restrict y, ulong n)
3  // n must be even
4  {
5      const ulong nh = n/2;
6      for (ulong k=0, k2=0;     k<nh;  k++, k2+=2)  y[k2] = x[k];
7      for (ulong k=nh, k2=n-1;  k<n;   k++, k2-=2)  y[k2] = x[k];
8  }
```

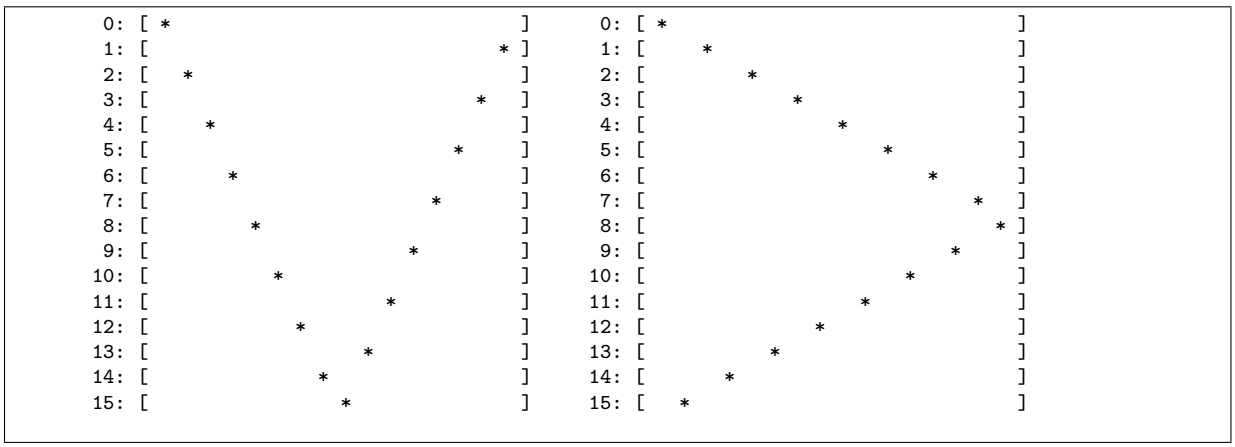

Figure 2.6-A: Permutation matrices of the reversed zip permutation (left) and its inverse (right).

The in-place version can, if the array length is a power of two, be implemented as

```
1  template <typename Type>
2  void zip_rev(Type *x, ulong n)
3  // n must be a power of two
4  {
5      const ulong nh = n/2;
6      reverse(x+nh, nh);
7      revbin_permute(x, nh); revbin_permute(x+nh, nh);
8      revbin_permute(x, n);
9  }
```

[1] For $n_r = n_c$ it of course is.

The inverse permutation [FXT: `unzip_rev()` in perm/ziprev.h] can be implemented as

```
1   template <typename Type>
2   void unzip_rev(const Type * restrict x, Type * restrict y, ulong n)
3   // n must be even
4   {
5       const ulong nh = n/2;
6       for (ulong k=0, k2=0;     k<nh;  k++, k2+=2)  y[k] = x[k2];
7       for (ulong k=nh, k2=n-1;  k<n;   k++, k2-=2)  y[k] = x[k2];
8   }
```

The in-place version is

```
1   template <typename Type>
2   void unzip_rev(Type *x, ulong n)
3   // n must be a power of two
4   {
5       const ulong nh = n/2;
6       revbin_permute(x, n);
7       revbin_permute(x, nh); revbin_permute(x+nh, nh);
8       reverse(x+nh, nh);
9   }
```

The given permutation is used in an algorithm where the cosine transform is computed using the Hartley transform, see section 24.11 on page 518.

We write Z and Z^{-1} for the zip permutation and its inverse, $\overline{Z}$ and $\overline{Z}^{-1}$ for the reversed zip permutation and its inverse, and R for the revbin permutation. Then the following relations hold:

$$Z = R\,Z^{-1}\,R = \overline{Z}\,Z^{-1}\,\overline{Z} \tag{2.6-1a}$$
$$\overline{Z} = R\,\overline{Z}^{-1}\,R = Z\,\overline{Z}^{-1}\,Z \tag{2.6-1b}$$
$$Z^{-1} = R\,Z\,R = \overline{Z}^{-1}\,Z\,\overline{Z}^{-1} \tag{2.6-1c}$$
$$\overline{Z}^{-1} = R\,\overline{Z}\,R = Z^{-1}\,\overline{Z}\,Z^{-1} \tag{2.6-1d}$$

2.7 The XOR permutation

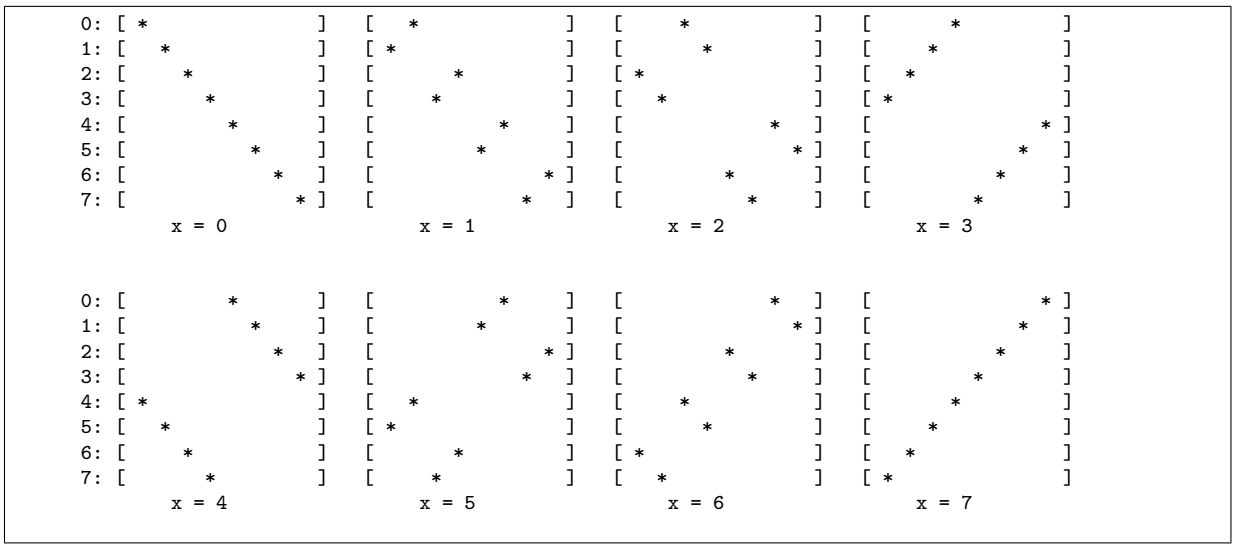

Figure 2.7-A: Permutation matrices of the XOR permutation for length 8 with parameter $x = 0 \ldots 7$. Compare to the table for the dyadic convolution shown in figure 22.7-A on page 464.

The *XOR permutation* may be explained most simply by its trivial implementation: [FXT: `xor_permute()` in perm/xorpermute.h]:

```
1   template <typename Type>
2   void xor_permute(Type *f, ulong n, ulong x)
3   {
4       if ( 0==x )  return;
5       for (ulong k=0; k<n; ++k)
6       {
7           ulong r = k^x;
8           if ( r>k )  swap2(f[r], f[k]);
9       }
10  }
```

The XOR permutation is evidently self-inverse. The array length n must be divisible by the smallest power of two that is greater than x: for example, n must be even if $x = 1$, n must be divisible by four if $x = 2$ or $x = 3$. With n a power of two and $x < n$ one is on the safe side.

The XOR permutation contains a few other permutations as important special cases (for simplicity assume that the array length n is a power of two): when the third argument x equals $n-1$ then the permutation is the reversion, with $x = 1$ neighboring even and odd indexed elements are swapped, with $x = n/2$ the upper and the lower half of the array are swapped.

One has

$$X_a\, X_b \;\; = \;\; X_b\, X_a \;=\; X_c \quad \text{where} \quad c = a \text{ XOR } b \tag{2.7-1}$$

For the special case $a = b$ the relation expresses the self-inverse property as X_0 is the identity. The XOR permutation often occurs in relations between other permutations where we will use the symbol X_x, the subscript denoting the third argument.

2.8 The Gray code permutation

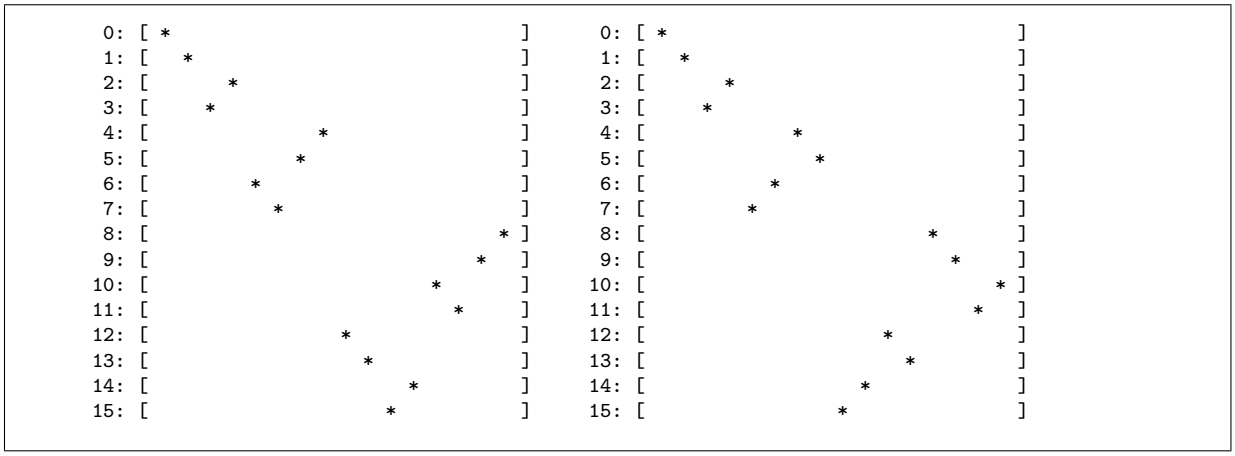

Figure 2.8-A: Permutation matrices of the Gray code permutation (left) and its inverse (right).

The *Gray code permutation* (or simply *Gray permutation*) reorders (length-2^n) arrays according to the binary Gray code described in section 1.16 on page 40. A routine for the permutation is [FXT: perm/graypermute.h]:

```
1   template <typename Type>
2   inline void gray_permute(const Type *f, Type * restrict g, ulong n)
3   // Put Gray permutation of f[] to g[], i.e. g[gray_code(k)] == f[k]
4   {
5       for (ulong k=0; k<n; ++k)  g[gray_code(k)] = f[k];
6   }
```

Its inverse is

```
1   template <typename Type>
2   inline void inverse_gray_permute(const Type *f, Type * restrict g, ulong n)
3   // Put inverse Gray permutation of f[] to g[], i.e. g[k] == f[gray_code(k)]
4   // (same as: g[inverse_gray_code(k)] == f[k])
5   {
6       for (ulong k=0; k<n; ++k)  g[k] = f[gray_code(k)];
7   }
```

We again use calls to `gray_code()` because they are cheaper than the computation of `inverse_gray_code()`.

We now give in-place versions of the above routines that offer very good performance. It is necessary to identify the cycle leaders (see section 2.11 on page 112) of the permutation and find an efficient way to generate them.

2.8.1 Cycles of the permutation

```
     0:  (  2,  3) #=2
     1:  (  4,  7,  5,  6) #=4
     2:  (  8, 15, 10, 12) #=4
     3:  (  9, 14, 11, 13) #=4
     4:  ( 16, 31, 21, 25, 17, 30, 20, 24) #=8
     5:  ( 18, 28, 23, 26, 19, 29, 22, 27) #=8
     6:  ( 32, 63, 42, 51, 34, 60, 40, 48) #=8
     7:  ( 33, 62, 43, 50, 35, 61, 41, 49) #=8
     8:  ( 36, 56, 47, 53, 38, 59, 45, 54) #=8
     9:  ( 37, 57, 46, 52, 39, 58, 44, 55) #=8
    10:  ( 64,127, 85,102, 68,120, 80, 96) #=8
    11:  ( 65,126, 84,103, 69,121, 81, 97) #=8
    12:  ( 66,124, 87,101, 70,123, 82, 99) #=8
    13:  ( 67,125, 86,100, 71,122, 83, 98) #=8
    14:  ( 72,112, 95,106, 76,119, 90,108) #=8
    15:  ( 73,113, 94,107, 77,118, 91,109) #=8
    16:  ( 74,115, 93,105, 78,116, 88,111) #=8
    17:  ( 75,114, 92,104, 79,117, 89,110) #=8
    126 elements in  18 nontrivial cycles.
    cycle lengths:   2 ...   8
  2 fixed points: [0. 1]
```

Figure 2.8-B: Cycles of the Gray code permutation of length 128.

It is instructive to study the complementary masks that occur for cycles (see section 2.11 on page 112) of different lengths. The cycles of the Gray code permutation for length 128 are shown in figure 2.8-B. No structure is immediately visible.

However, one can generate the cycle maxima as follows: for each range $2^k \ldots 2^{k+1} - 1$ generate a bit-mask z that is obtained from the $k+1$ leftmost bits of the infinite word that has bits set at positions 0, 1, 2, 4, 8, ..., 2^i, ...:

```
[11101000100000001000000000000001000 ...]
```

An example: for $k = 6$ we have $z =$`[1110100]`. Then take v to be $k+1$ leftmost bits of the complement, $v =$`[0001011]` in our example. Now the set of words $c = z + s$ where s is a subset of v contains exactly one element of each cycle in the range $2^k \ldots 2^{k+1} = 64 \ldots 127$:

```
117 =  .111.1.1
118 =  .111.11.
119 =  .111.111
124 =  .11111..
125 =  .11111.1
126 =  .111111.
127 =  .1111111
116 =  .111.1..
  maxima = z XOR subsets(v)  where
  z =  .111.1..
  v =  ....1.11
```

The words obtained are actually the cycle maxima. The list can be generated with the program [FXT: perm/permgray-leaders-demo.cc] which uses [FXT: `class gray_cycle_leaders` in comb/gray-cycle-leaders.h]:

```
class gray_cycle_leaders
// Generate cycle leaders for Gray code permutation
// where highest bit is at position ldn.
{
public:
    bit_subset b_;
    ulong za_;  // mask for cycle maxima
    ulong zi_;  // mask for cycle minima
    ulong len_; // cycle length
    ulong num_; // number of cycles

public:

    gray_cycle_leaders(ulong ldn)  // 0<=ldn<BITS_PER_LONG
        : b_(0)
    {
        init(ldn);
    }

    ~gray_cycle_leaders() {;}

    void init(ulong ldn)
    {
        za_ = 1;
        ulong cz = 0;   // ~z
        len_ = 1;
        num_ = 1;
        for (ulong ldm=1; ldm<=ldn; ++ldm)
        {
            za_ <<= 1;
            cz <<= 1;
            if ( is_pow_of_2(ldm) )
            {
                ++za_;
                len_ <<= 1;
            }
            else
            {
                ++cz;
                num_ <<= 1;
            }
        }

        zi_ = 1UL << ldn;

        b_.first(cz);
    }

    ulong current_max()  const  { return b_.current() | za_; }
    ulong current_min()  const  { return b_.current() | zi_; }

    bool next()
    {
        return ( 0!=b_.next() );
    }

    ulong num_cycles()  const  { return num_; }
    ulong cycle_length()  const  { return len_; }
};
```

The implementation uses the bit-subset class described in section 1.27 on page 70.

2.8.2 In-place routines

For the in-place versions of the permutation routines are obtained by inlining the generation of the cycle leaders. The forward version is [FXT: perm/graypermute.h]:

```
template <typename Type>
void gray_permute(Type *f, ulong n)
{
```

```
4       ulong z = 1; // mask for cycle maxima
5       ulong v = 0; // ~z
6       ulong cl = 1;  // cycle length
7       for (ulong ldm=1, m=2;  m<n;  ++ldm, m<<=1)
8       {
9           z <<= 1;
10          v <<= 1;
11          if ( is_pow_of_2(ldm) )
12          {
13              ++z;
14              cl <<= 1;
15          }
16          else  ++v;
17
18          bit_subset b(v);
19          do
20          {
21              // --- do cycle: ---
22              ulong i = z | b.next();  // start of cycle
23              Type t = f[i];           // save start value
24              ulong g = gray_code(i);  // next in cycle
25              for (ulong k=cl-1; k!=0; --k)
26              {
27                  Type tt = f[g];
28                  f[g] = t;
29                  t = tt;
30                  g = gray_code(g);
31              }
32              f[g] = t;
33              // --- end (do cycle) ---
34          }
35          while ( b.current() );
36      }
37  }
```

The function `is_pow_of_2()` is described in section 1.7 on page 18. The inverse routine differs only in the block that processes the cycles:

```
1   template <typename Type>
2   void inverse_gray_permute(Type *f, ulong n)
3   {
4       [--snip--]
5               // --- do cycle: ---
6               ulong i = z | b.next();  // start of cycle
7               Type t = f[i];           // save start value
8               ulong g = gray_code(i);  // next in cycle
9               for (ulong k=cl-1; k!=0; --k)
10              {
11                  f[i] = f[g];
12                  i = g;
13                  g = gray_code(i);
14              }
15              f[i] = t;
16              // --- end (do cycle) ---
17      [--snip--]
18  }
```

2.8.3 Performance of the routines

How fast is the Gray code permutation? We use the convention that the speed of the trivial (and completely cache-friendly, therefore running at memory bandwidth) `reverse()` is 1.0, our hereby declared time unit for comparison [FXT: perm/reverse.h]. A little benchmark gives, for large (16 MB) arrays:

```
arg 1: 21 == ldn  [Using 2**ldn elements]  default=21
arg 2: 10 == rep  [Number of repetitions]  default=10
Memsize = 16384 kiloByte  ==  2097152 doubles

                reverse(f,n);    dt=  0.0103524    MB/s= 1546   rel=          1
         revbin_permute(f,n);    dt=  0.0674235    MB/s=  237   rel=    6.51282
        revbin_permute0(f,n);    dt=   0.061507    MB/s=  260   rel=    5.94131
           gray_permute(f,n);    dt=  0.0155019    MB/s= 1032   rel=    1.49742
   inverse_gray_permute(f,n);    dt=  0.0150641    MB/s= 1062   rel=    1.45512
```

```
                reverse(f,n);    dt=  0.0104008    MB/s= 1538   rel=    1.00467
```

We timed `reverse()` twice to get an impression how much we can trust the observed numbers.

While the revbin permutation takes about 6 units (due to its memory access pattern that is very problematic with respect to cache usage) the Gray code permutation needs only 1.50 units. The difference gets bigger for machines with relatively (to the CPU) slow memory. The Gray code permutation can be used to speed up fast transforms of large lengths a power of two, notably the Walsh transform, see chapter 22 on page 447.

The bandwidth of the `reverse()` is about 1500 MB/sec which should be compared to the output of a memory testing program, revealing that it actually runs at about the bandwidth of copying via a simple loop using pointers to `double`s:

```
avg:   16777216  [ 0]"memcpy"                   2522.084 MB/s
avg:   16777216  [ 1]"char *"                    471.873 MB/s
avg:   16777216  [ 2]"short *"                   711.853 MB/s
avg:   16777216  [ 3]"int *"                     956.682 MB/s
avg:   16777216  [ 4]"long *"                   1514.360 MB/s
avg:   16777216  [ 5]"long * (4x unrolled)"    1330.786 MB/s
avg:   16777216  [ 6]"int64 *"                  1329.902 MB/s
avg:   16777216  [ 7]"double *"                1329.507 MB/s  // <--=
avg:   16777216  [ 8]"double * (4x unrolled)"  1325.437 MB/s
```

The relative speeds are quite different for small arrays. Using a size of 16 kB (2048 `double`s) we obtain

```
arg 1: 11 == ldn  [Using 2**ldn elements]  default=21
arg 2: 100000 == rep  [Number of repetitions]  default=512
Memsize = 16 kiloByte  ==  2048 doubles
                reverse(f,n);    dt=1.88726e-06    MB/s= 8279   rel=          1
          revbin_permute(f,n);    dt=3.22166e-06    MB/s= 4850   rel=    1.70706
         revbin_permute0(f,n);    dt=2.69212e-06    MB/s= 5804   rel=    1.42647
            gray_permute(f,n);    dt=4.75155e-06    MB/s= 3288   rel=    2.51769
    inverse_gray_permute(f,n);    dt=3.69237e-06    MB/s= 4232   rel=    1.95647
                reverse(f,n);    dt=1.88833e-06    MB/s= 8275   rel=    1.00057
```

due to the small size, the cache problems are gone.

The memory benchmark gives for that size

```
avg:      16384  [ 0]"memcpy"                   3290.353 MB/s
avg:      16384  [ 1]"char *"                    572.922 MB/s
avg:      16384  [ 2]"short *"                   973.552 MB/s
avg:      16384  [ 3]"int *"                    1495.920 MB/s
avg:      16384  [ 4]"long *"                   3560.506 MB/s
avg:      16384  [ 5]"long * (4x unrolled)"    3220.792 MB/s
avg:      16384  [ 6]"int64 *"                  2498.137 MB/s
avg:      16384  [ 7]"double *"                2498.285 MB/s  // <--=
avg:      16384  [ 8]"double * (4x unrolled)"  3219.784 MB/s
```

2.9 The reversed Gray code permutation

If the length-n array is permuted in the way the upper half of the length-$2n$ array would be permuted by `gray_permute()` then all cycles are of the same length. The resulting permutation is equivalent to the *reversed Gray code permutation*:

```
1  template <typename Type>
2  inline void gray_rev_permute(const Type *f, Type * restrict g, ulong n)
3  // gray_rev_permute() =^=
4  //  { reverse(); gray_permute(); }
5  {
6      for (ulong k=0, m=n-1;  k<n;  ++k, --m)  g[gray_code(m)] = f[k];
7  }
```

The routine, its inverse and in-place versions are given in [FXT: perm/grayrevpermute.h].

All cycles have the same length, `gray_rev_permute(f, 64)` gives:

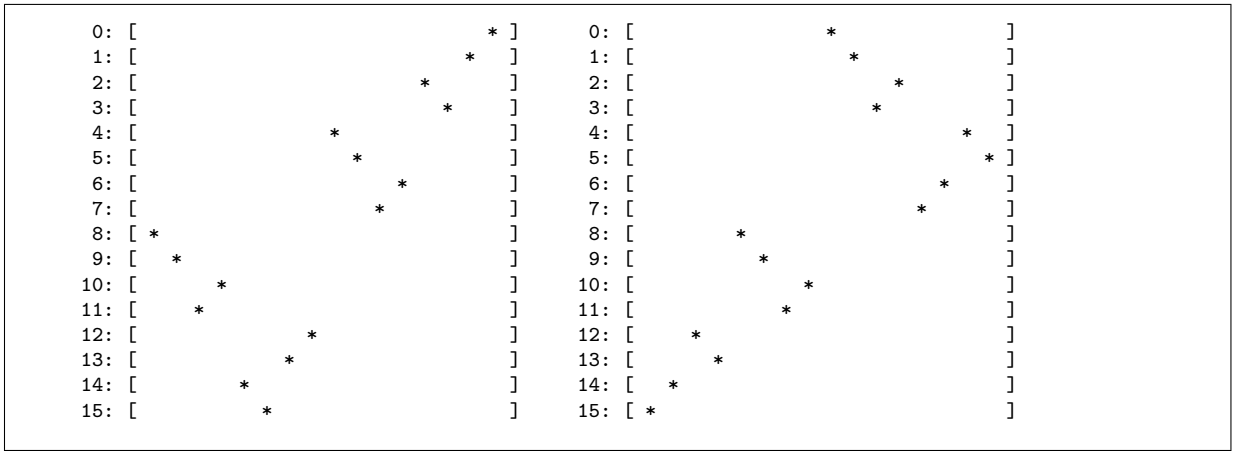

Figure 2.9-A: Permutation matrices of the reversed Gray code permutation (left) and its inverse (right).

```
0:  (  0, 63, 21, 38,  4, 56, 16, 32) #=8
1:  (  1, 62, 20, 39,  5, 57, 17, 33) #=8
2:  (  2, 60, 23, 37,  6, 59, 18, 35) #=8
3:  (  3, 61, 22, 36,  7, 58, 19, 34) #=8
4:  (  8, 48, 31, 42, 12, 55, 26, 44) #=8
5:  (  9, 49, 30, 43, 13, 54, 27, 45) #=8
6:  ( 10, 51, 29, 41, 14, 52, 24, 47) #=8
7:  ( 11, 50, 28, 40, 15, 53, 25, 46) #=8
64 elements in   8 nontrivial cycles.
cycle length is == 8
No fixed points.
```

If 64 is added to the indices then the cycles in the upper half of the array as in `gray_permute(f, 128)` are reproduced (by construction).

Let G denote the Gray code permutation, $\overline{G}$ the reversed Gray code permutation. Symbolically one can write

$$G(n) = \{\dots, \overline{G}(n/8), \overline{G}(n/4), \overline{G}(n/2)\} \tag{2.9-1a}$$

$$G^{-1}(n) = \{\dots, \overline{G}^{-1}(n/8), \overline{G}^{-1}(n/4), \overline{G}^{-1}(n/2)\} \tag{2.9-1b}$$

Now let r be the reversion and h the permutation that swaps the upper and the lower half of an array. Then

$$\overline{G} = G\,r = h\,G \tag{2.9-2a}$$

$$\overline{G}^{-1} = r\,G^{-1} \tag{2.9-2b}$$

$$\overline{G}^{-1}\,G = G^{-1}\,\overline{G} = r = X_{n-1} \tag{2.9-2c}$$

$$G\,\overline{G}^{-1} = \overline{G}\,G^{-1} = h = X_{n/2} \tag{2.9-2d}$$

Throughout it is assumed that the array length n is a power of two.

2.10 Decomposing permutations *

In this section we will see some algorithms that use a certain type of decomposition (factorization in term of matrices) of some of the permutations we have studied so far. The resulting algorithms involve proportional $n \cdot \log(n)$ computations for length-n arrays. This might seem to render the schemes worthless as one can always obtain a permutation with work proportional to n. There are, however, situations where one can use the algorithms advantageously. Firstly, with bit manipulations, the whole binary words are modified in one or a few statements: the corresponding algorithms are therefore only proportional $\log(n)$. Secondly, when the algorithm can be used implicitly in order to integrate the permutation in a fast transform. The work can sometimes be reduced to zero in that case.

2.10.1 Array reversal

The most simple example might be the reversion via

```
1   void perm1(ulong *f, ulong n)
2   // From shorter to longer sub-arrays.
3   {
4       for (ulong k=2; k<=n; k*=2)
5       {
6           for (ulong j=0; j<n; j+=k)  func(f+j, k);
7       }
8   }
```

where `func()` swaps the upper and lower half of an array:

```
1   void func(ulong *f, ulong ldn)  { swap(f, f+n/2, n/2); }
```

This idea has been exploited in section 1.14 on page 34 in order to obtain a bit-reversal routine. The reversal is a self-inverse operation. Therefore one can alternatively execute the steps in reverse order and still get the same permutation:

```
1   void perm2(ulong *f, ulong n)
2   // From longer to shorter sub-arrays.
3   {
4       for (ulong k=n; k>=2; k/=2) // k == n, n/2, n/4, ... , 4, 2
5       {
6           for (ulong j=0; j<n; j+=k)  func(f+j, k);
7       }
8   }
```

Note that `func()` in turn can be obtained via

```
reverse(f, n);
reverse(f, n/2);  reverse(f+n/2, n/2);
```

or the same statements in reversed order.

2.10.2 Gray code permutation

Let us try a less obvious example, use `perm1()` with `func()` defined to swap the third and fourth quarter:

```
swap(f+n/2, f+n/2+n/4, n/4);  // quarters: [0,1,2,3]-->[0,1,3,2]
```

The resulting permutation is the Gray permutation. The other way round (using `perm2()`) one obtains the inverse Gray permutation.

Using `func()` defined as

```
reverse(f+n/2, n/2);
```

one gets the Gray permutation through `perm2()`, its inverse through `perm1()`. This idea has been used in the core-routine for the sequency-ordered Walsh transform described in section 22.9 on page 466. The work for the Gray permutation has been completely vaporized there.

Note that the routine that swaps the halves of the upper half array could be obtained as either of

```
inverse_gray_permute(f, n);
gray_permute(f, n/2);     gray_permute(f+n/2, n/2);
```

or

```
inverse_gray_permute(f, n/2);  inverse_gray_permute(f+n/2, n/2);
gray_permute(f, n);
```

Similarly, the routine that reverses the upper half can be obtained as either of

```
gray_permute(f, n/2);     gray_permute(f+n/2, n/2);
inverse_gray_permute(f, n);
```

or

```
gray_permute(f, n);
inverse_gray_permute(f, n/2);  inverse_gray_permute(f+n/2, n/2);
```

The corresponding routines to Gray-permute the bits of a binary word are given in [FXT: bits/bitgraypermute.h].

2.10.3 Zip and revbin permutation

Using `func()` defined to swap the second and third quarter

```
swap(f+n/4, f+n/2, n/4);  // quarters: [0,1,2,3]-->[0,2,1,3]
```

Then with `perm2()` one gets the zip permutation, `perm1()` gives the inverse. This idea has been used for the bit-wise zip shown in section 1.15 on page 39.

Using `func()` to cycle the second, third and fourth quarter:

```
// quarters: [0,1,2,3]-->[0,2,3,1]
```

which was obtained using

```
zip_rev(f, n);
unzip_rev(f, n/2);  unzip_rev(f+n/2, n/2);
```

both the reversed zip permutation and its inverse can be computed in the now hopefully obvious way.

The revbin permutation can be generated through the zip permutation or its inverse. However, the zip permutation is the more complicate one, so absorbing the revbin permutation into fast transforms does not seem to be easy. The other way round it makes more sense:

```
revbin_permute(f, n/2);  revbin_permute(f+n/2, n/2);
revbin_permute(f, n);
```

Is a convenient (though not the most effective) way to compute the zip permutation.

Clearly, the idea presented here is in analogy with the decomposition of linear transforms. Finally the permutations are (very simple forms of) linear transforms. See also section 1.33 on page 88.

2.11 General permutations and their operations

So far we treated special permutations that are used in certain algorithms. Now we study permutations in general with the operations such as composition and inversion on them.

2.11.1 Basic definitions and operations

A straightforward way to represent a permutation is to consider the array of indices that for the original (unpermuted) data would be the length-n *canonical sequence* 0, 1, 2, ..., $n-1$. The mentioned trivial sequence represents the 'do-nothing' permutation or *identity*. The concept is best described by the routine that *applies* a given permutation x on an array of data f: after the routine has finished the array g will contain the elements of f reordered according to x [FXT: `apply_permutation()` in perm/permapply.h]:

```
1  template <typename Type>
2  void apply_permutation(const ulong *x, const Type *f, Type * restrict g, ulong n)
3  // Apply the permutation x[] on f[]
4  // i.e. set g[k] <-- f[x[k]]  \forall k
5  {
6      for (ulong k=0; k<n; ++k)  g[k] = f[x[k]];
7  }
```

An example using strings (arrays of characters): the permutation represented by

$x =$`[ 7 6 3 2 5 1 0 4 ]`

and the input data

$f =$`[ A B a d C a f e ]` would produce

$g =$`[ e f d a a B A C ]`

Routines that test various properties of permutations are given in [FXT: perm/permq.h]. To check whether a given permutation is the identity is trivial:

```
bool is_identity(const ulong *f, ulong n)
// Return whether f[] is the identical permutation,
// i.e. whether f[k]==k for all k= 0...n-1
{
    for (ulong k=0; k<n; ++k)  if ( f[k] != k )  return false;
    return true;
}
```

A fixed point of a permutation is an index where the element is not moved:

```
ulong count_fixed_points(const ulong *f, ulong n)
// Return number of fixed points in f[]
{
    ulong ct = 0;
    for (ulong k=0; k<n; ++k)  if ( f[k] == k )  ++ct;
    return ct;
}
```

A *derangement* is a permutation that has no fixed points. To check whether a permutation is a derangement of identity use:

```
bool is_derangement(const ulong *f, ulong n)
// Return whether f[] is a derangement of identity,
// i.e. whether f[k]!=k for all k
{
    for (ulong k=0; k<n; ++k)  if ( f[k] == k )  return false;
    return true;
}
```

Whether two arrays are mutual derangements can be determined by:

```
bool is_derangement(const ulong *f, const ulong *g, ulong n)
// Return whether f[] is a derangement of g[],
// i.e. whether f[k]!=g[k] for all k
{
    for (ulong k=0; k<n; ++k)  if ( f[k] == g[k] )  return false;
    return true;
}
```

To check whether a given array really describes a valid permutation one has to verify that each index in the valid range appears exactly once. The `bitarray` class described in section 4.6 on page 160 allows us to do the job without modification of the input:

```
bool
is_valid_permutation(const ulong *f, ulong n, bitarray *bp/*=0*/)
// Return whether all values 0...n-1 appear exactly once,
// i.e. whether f represents a permutation of [0,1,...,n-1].
{
    // check whether any element is out of range:
    for (ulong k=0; k<n; ++k)  if ( f[k]>=n )  return false;

    // check whether values are unique:
    bitarray *tp = bp;
    if ( 0==bp )  tp = new bitarray(n);  // tags
    tp->clear_all();

    ulong k;
    for (k=0; k<n; ++k)
    {
        if ( tp->test_set(f[k]) )  break;
    }

    if ( 0==bp )  delete tp;

    return  (k==n);
}
```

We note two rather trivial operations for permutation, computing the *complement* [FXT: perm/permcomplement.h]

```
inline void make_complement(const ulong *f, ulong *g, ulong n)
// Set (as permutation) g to the complement of f.
// Can have f==g.
```

```
4   {
5       for (ulong k=0; k<n; ++k)  g[k] = n - 1 - f[k];
6   }
```

and computing the reversal [FXT: perm/reverse.h]

```
1   template <typename Type>
2   inline void reverse(Type *f, ulong n)
3   // Reverse order of array f.
4   {
5       for (ulong k=0, i=n-1;  k<i;  ++k, --i)  swap2(f[k], f[i]);
6   }
```

2.11.2 Compositions of permutations

One can apply several permutations to an array, one by one. The resulting permutation is called the *composition* of the applied permutations. The routines are given in [FXT: perm/permq.cc]. As an example, the check whether some permutation g is equal to f applied twice, or f *squared*, use:

```
1   bool is_square(const ulong *f, const ulong *g, ulong n)
2   // Return whether f * f == g  as a permutation
3   {
4       for (ulong k=0; k<n; ++k)  if ( g[k] != f[f[k]] )  return false;
5       return true;
6   }
```

Note that in general $f \cdot g \neq g \cdot f$ for $f \neq g$, the operation of composition is not commutative.

A permutation f is said to be the *inverse* of another permutation g if it undoes its effect, that is $f \cdot g = \mathrm{id}$:

```
1   bool is_inverse(const ulong *f, const ulong *g, ulong n)
2   // Return whether f[] is the inverse of g[]
3   {
4       for (ulong k=0; k<n; ++k)  if ( f[g[k]] != k )  return false;
5       return true;
6   }
```

One has $g \cdot f = f \cdot g = \mathrm{id}$, in a group the left-inverse is equal to the right-inverse and we can simply call g 'the inverse' of f.

A permutation that is its own inverse (like the revbin permutation) is called an *involution*. Checking that is easy:

```
1   bool is_involution(const ulong *f, ulong n)
2   // Return whether max cycle length is <= 2,
3   // i.e. whether f * f = id.
4   {
5       for (ulong k=0; k<n; ++k)  if ( f[f[k]] != k )  return false;
6       return true;
7   }
```

The following routine computes the inverse of a given permutation [FXT: perm/perminvert.cc]:

```
1   void make_inverse(const ulong *f, ulong * restrict g, ulong n)
2   // Set (as permutation) g to the inverse of f
3   {
4       for (ulong k=0; k<n; ++k)  g[f[k]] = k;
5   }
```

2.11.3 Representation as disjoint cycles

If one wants to do the operation in-place a little bit of thought is required. The idea underlying all subsequent routines working in-place is that every permutation entirely consists of disjoint cycles. A *cycle* of a permutation is a subset of the indices that is rotated (by one) by the permutation. The term *disjoint* means that the cycles do not 'cross' each other. While this observation is pretty trivial it allows us to do many operations by following the cycles of the permutation, one by one, and doing the necessary operation on each of them. As an example consider the following permutation of an array originally consisting of the (canonical) sequence 0, 1, ..., 15. Extra spaces are inserted for readability:

```
[ 0, 1, 3, 2,   7, 6, 4, 5,   15, 14, 12, 13,   8, 9, 11, 10 ]
```

There are two fixed points (0 and 1) and these cycles:

```
( 2 <--  3 )
( 4 <--  7 <--  5 <--  6 )
( 8 <-- 15 <-- 10 <-- 12 )
( 9 <-- 14 <-- 11 <-- 13 )
```

The cycles do 'wrap around', for example, the initial 4 of the second cycle goes to position 6, the last element of the second cycle.

Note that the inverse permutation could formally be described by reversing every arrow in each cycle:

```
( 2 -->  3 )
( 4 -->  7 -->  5 -->  6 )
( 8 --> 15 --> 10 --> 12 )
( 9 --> 14 --> 11 --> 13 )
```

Equivalently, one can reverse the order of the elements in each cycle:

```
( 3 <--  2 )
( 6 <--  5 <--  7 <--  4 )
(12 <-- 10 <-- 15 <--  8 )
(13 <-- 11 <-- 14 <--  9 )
```

If we begin each cycle with its smallest element the inverse permutation is written as

```
( 2 <--  3 )
( 4 <--  6 <--  5 <--  7 )
( 8 <-- 12 <-- 10 <-- 15 )
( 9 <-- 13 <-- 11 <-- 14 )
```

This form is obtained by reversing all elements except the first in each cycle of the (forward) permutation. The last three sets of cycles all describe the same permutation:

```
[ 0, 1, 3, 2,   6, 7, 5, 4,   12, 13, 15, 14,   10, 11, 9, 8 ]
```

The cycles above were printed with [FXT: `print_cycles()` in perm/printcycles.cc]

```
1   ulong print_cycles(const ulong *f, ulong n, bitarray *bp=0)
2   // Print the cycles of the permutation.
3   // Return number of fixed points.
4   {
5       bitarray *tp = bp;
6       if ( 0==bp )  tp = new bitarray(n);  // tags
7       tp->clear_all();
8
9       ulong ct = 0;  // # of fixed points
10      for (ulong k=0; k<n; ++k)
11      {
12          if ( tp->test_clear(k) )  continue;  // already processed
13          tp->set(k);
14
15          // follow a cycle:
16          ulong i = k;
17          ulong g = f[i];  // next index
18          if ( g==i )  // fixed point ?
19          {
20              ++ct;
21              continue;
22          }
23
24          cout << "(" << setw(3) << i;
25          while ( 0==(tp->test_set(g)) )
26          {
27              cout << " <-- " << setw(3) << g;
28              g = f[g];
29          }
30          cout << " )" << endl;
31      }
32
33      if ( 0==bp )  delete tp;
34
35      return  ct;
36  }
```

The bit-array (see section 4.6 on page 160 for the implementation) is used to keep track of the elements already processed.

A utility class to compute the decomposition of a permutation into cycles is [FXT: `class cycles` in perm/cycles.h]. A program that shows its usage is [FXT: perm/cycles-demo.cc], it prints the cycles

```
Using: gray_permute(y, n)
Computing cycles:
  0:  (  2,  3) #=2
  1:  (  4,  7,  5,  6) #=4
  2:  (  8, 15, 10, 12) #=4
  3:  (  9, 14, 11, 13) #=4
  14 elements in   4 nontrivial cycles.
 cycle lengths:   2 ...   4
number of fixed points = 2
```

and code for a permutation of given size

```
1   template <typename Type>
2   inline void foo_perm_16(Type *f)
3   // unrolled version for length 16
4   {
5     swap2(f[2], f[3]);
6     { Type t=f[4]; f[4]=f[7]; f[7]=f[5]; f[5]=f[6]; f[6]=t; }
7     { Type t=f[8]; f[8]=f[15]; f[15]=f[10]; f[10]=f[12]; f[12]=t; }
8     { Type t=f[9]; f[9]=f[14]; f[14]=f[11]; f[11]=f[13]; f[13]=t; }
9   }
```

2.11.4 Cyclic permutations

A permutation consisting of exactly one cycle is called *cyclic*. Whether a given permutation has this property can be tested with [FXT: `is_cyclic()` in perm/permq.cc]:

```
1   bool
2   is_cyclic(const ulong *f, ulong n)
3   // Return whether permutation is exactly one cycle.
4   {
5       if ( n<=1 )  return true;
6       ulong k = 0,  e = 0;
7       do  { e=f[e]; ++k; }  while ( e!=0 );
8       return  (k==n);
9   }
```

The method used is to follow the cycle starting at position zero and counting how long it is. The permutation is cyclic exactly if the length found equals the array length. There are $(n-1)!$ cyclic permutations of n elements.

2.11.5 Sign and parity of a permutation

Every permutation can be written as a composition of *transpositions* (cycles of length two). This composition is not unique, but its number modulo two is unique. The *sign* of a permutation is defined to be $+1$ the number is even and -1 if the number is odd. The minimal number of transpositions whose composition give a cycle of length l is $l-1$. So the minimal number of transpositions for a permutation consisting of k cycles where the length of the j-th cycle is l_j equals $\sum_{j=1}^{k} l_j - 1 = (\sum_{j=1}^{k} l_j) - k$. The sign corresponds to the homomorphic mapping into the group of the elements $+1$ and -1 with multiplication as group operation. If we count the transpositions modulo two (corresponding to the mapping into the additive group modulo two) we obtain what may be called the *parity* of a permutation.

2.11.6 Inverse and square, in-place

For the computation of the inverse we have to reverse each cycle [FXT: perm/perminvert.cc]:

```
1   void make_inverse(ulong *f, ulong n, bitarray *bp/*=0*/)
2   // Set (as permutation) f to its own inverse.
3   // In-place version.
4   {
5       bitarray *tp = bp;
6       if ( 0==bp )  tp = new bitarray(n);  // tags
7       tp->clear_all();
8
9       for (ulong k=0; k<n; ++k)
10      {
11          if ( tp->test_clear(k) )  continue;  // already processed
12          tp->set(k);
13
14          // invert a cycle:
15          ulong i = k;
16          ulong g = f[i];  // next index
17          while ( 0==(tp->test_set(g)) )
18          {
19              ulong t = f[g];
20              f[g] = i;
21              i = g;
22              g = t;
23          }
24          f[g] = i;
25      }
26
27      if ( 0==bp )  delete tp;
28  }
```

The extra array of tag bits can be avoided by using the highest bit of each word as tag bit. The scheme would fail if any word of the permutation array had the highest bit set. However, on byte-addressable machines such an array will not fit into memory at all (for word sizes of 16 or more bits). To keep the code similar to the version using the bit-array we define

```
1   static const ulong s1 = 1UL << (BITS_PER_LONG - 1);  // highest bit is tag bit
2   static const ulong s0 = ~s1;  // all bits but tag bit
3
4   static inline void SET(ulong *f, ulong k)  { f[k&s0] |= s1; }
5   static inline void CLEAR(ulong *f, ulong k)  { f[k&s0] &= s0; }
6   static inline bool TEST(ulong *f, ulong k)  { return (0!=(f[k&s0]&s1)); }
```

Note that we have to mask out the tag-bit when using the value '`k`' as index. The routine can then be implemented as

```
1   void
2   make_inverse(ulong *f, ulong n)
3   // Set (as permutation) f to its own inverse.
4   // In-place version using highest bits of array as tag-bits.
5   {
6       for (ulong k=0; k<n; ++k)
7       {
8           if ( TEST(f, k) ) { CLEAR(f, k);  continue; }  // already processed
9           SET(f, k);
10
11          // invert a cycle:
12          ulong i = k;
13          ulong g = f[i];  // next index
14          while ( 0==TEST(f, g) )
15          {
16              ulong t = f[g];
17              f[g] = i;
18              SET(f, g);
19              i = g;
20              g = t;
21          }
22          f[g] = i;
23
24          CLEAR(f, k);  // leave no tag bits set
25      }
26  }
```

The extra `CLEAR()` statement at the end removes the tag-bit of the cycle minima. Its effect is that no tag-bits are set after routine has finished. The routine has about the same performance as the bit-array version. For the routine [FXT: perm/permcompose.cc]

```
1   void make_square(const ulong *f, ulong * restrict g, ulong n)
2   // Set (as permutation) g = f * f
3   {
4       for (ulong k=0; k<n; ++k)  g[k] = f[f[k]];
5   }
```

we obtain the following in-place version:

```
1   void make_square(ulong *f, ulong n, bitarray *bp/*=0*/)
2   // Set (as permutation) f = f * f
3   // In-place version.
4   {
5       bitarray *tp = bp;
6       if ( 0==bp )  tp = new bitarray(n);  // tags
7       tp->clear_all();
8
9       for (ulong k=0; k<n; ++k)
10      {
11          if ( tp->test_clear(k) )  continue;  // already processed
12          tp->set(k);
13
14          // square a cycle:
15          ulong i = k;
16          ulong t = f[i];  // save
17          ulong g = f[i];  // next index
18          while ( 0==(tp->test_set(g)) )
19          {
20              f[i] = f[g];
21              i = g;
22              g = f[g];
23          }
24          f[i] = t;
25      }
26
27      if ( 0==bp )  delete tp;
28  }
```

2.11.7 Powers of a permutation

The e-th power of a permutation f is computed (and returned in g) by a version of the binary exponentiation algorithm described in section 27.6 on page 555 [FXT: perm/permcompose.cc]:

```
1   void
2   power(const ulong *f, ulong * restrict g, ulong n, long e,
3         ulong * restrict t/*=0*/)
4   // Set (as permutation) g = f ** e
5   {
6       if ( e==0 )
7       {
8           for (ulong k=0; k<n; ++k)  g[k] = k;
9           return;
10      }
11
12      if ( e==1 )
13      {
14          copy(f, g, n);
15          return;
16      }
17
18      if ( e==-1 )
19      {
20          make_inverse(f, g, n);
21          return;
22      }
23
24      // here:  abs(e) > 1
25      ulong x = e>0 ? e : -e;
26
27      if ( is_pow_of_2(x) )  // special case x==2^n
28      {
29          make_square(f, g, n);
30          while ( x>2 )  { make_square(g, n);  x /= 2; }
31      }
```

```
32      else
33      {
34          ulong *tt = t;
35          if ( 0==t )  { tt = new ulong[n]; }
36          copy(f, tt, n);
37
38          int firstq = 1;
39          while ( 1 )
40          {
41              if ( x&1 )  // odd
42              {
43                  if ( firstq )  // avoid multiplication by 1
44                  {
45                      copy(tt, g, n);
46                      firstq = 0;
47                  }
48                  else  compose(tt, g, n);
49
50                  if ( x==1 )   goto dort;
51              }
52
53              make_square(tt, n);
54              x /= 2;
55          }
56
57      dort:
58          if ( 0==t )  delete [] tt;
59      }
60
61      if ( e<0 )  make_inverse(g, n);
62  }
```

2.11.8 Applying permutations to data, in-place

The in-place analogue for the routine

```
1   template <typename Type>
2   void apply_permutation(const ulong *x, const Type *f, Type * restrict g, ulong n)
3   // Apply the permutation x[] to the array f[]
4   // i.e. set g[k] <-- f[x[k]]  \forall k
5   {
6       for (ulong k=0; k<n; ++k)  g[k] = f[x[k]];
7   }
```

is [FXT: perm/permapply.h]:

```
1   template <typename Type>
2   void apply_permutation(const ulong *x, Type * restrict f, ulong n, bitarray *bp=0)
3   // Apply the permutation x[] to the array f[]
4   // i.e. set f[k] <-- f[x[k]]  \forall k
5   // In-place version.
6   {
7       bitarray *tp = bp;
8       if ( 0==bp )  tp = new bitarray(n);  // tags
9       tp->clear_all();
10
11      for (ulong k=0; k<n; ++k)
12      {
13          if ( tp->test_clear(k) )  continue;  // already processed
14          tp->set(k);
15
16          // --- do cycle: ---
17          ulong i = k;  // start of cycle
18          Type t = f[i];
19          ulong g = x[i];
20          while ( 0==(tp->test_set(g)) )  // cf. inverse_gray_permute()
21          {
22              f[i] = f[g];
23              i = g;
24              g = x[i];
25          }
26          f[i] = t;
27          // --- end (do cycle) ---
```

```
28      }
29
30      if ( 0==bp )  delete tp;
31  }
```

To apply the inverse of a permutation without actually inverting the permutation itself use

```
1   template <typename Type>
2   void apply_inverse_permutation(const ulong *x, const Type *f, Type * restrict g, ulong n)
3   // Apply the inverse permutation of x[] to the array f[],
4   // i.e. set g[x[k]] <-- f[k]  \forall k
5   {
6       for (ulong k=0; k<n; ++k)  g[x[k]] = f[k];
7   }
```

The in-place version is

```
1   template <typename Type>
2   void apply_inverse_permutation(const ulong *x, Type * restrict f, ulong n, bitarray *bp=0)
3   // Apply the inverse permutation of x[] to the array f[]
4   // i.e. set f[x[k]] <-- f[k]  \forall k
5   // In-place version.
6   {
7       bitarray *tp = bp;
8       if ( 0==bp )  tp = new bitarray(n);  // tags
9       tp->clear_all();
10
11      for (ulong k=0; k<n; ++k)
12      {
13          if ( tp->test_clear(k) )  continue;  // already processed
14          tp->set(k);
15
16          // --- do cycle: ---
17          ulong i = k;  // start of cycle
18          Type t = f[i];
19          ulong g = x[i];
20          while ( 0==(tp->test_set(g)) )  // cf. gray_permute()
21          {
22              Type tt = f[g];
23              f[g] = t;
24              t = tt;
25              g = x[g];
26          }
27          f[g] = t;
28          // --- end (do cycle) ---
29      }
30
31      if ( 0==bp )  delete tp;
32  }
```

When a permutation of the set $S := \{0, 1, \ldots, n-1\}$ is given as a function X (where $X(S) = S$) the permutation can be applied to an array f via [FXT: `apply_permutation()` in perm/permapplyfunc.h]:

```
1   template <typename Type>
2   void apply_permutation(ulong (*x)(ulong), const Type *f, Type * restrict g, ulong n)
3   // Set g[k] <-- f[x(k)]  \forall k
4   // Must have:  0<=x(k)<n \forall k
5   {
6       for (ulong k=0; k<n; ++k)  g[k] = f[x(k)];
7   }
```

For example, the statement `apply_permutation(inverse_gray_code, f, g, n)` is equivalent to `gray_permute(f, g, n)`. The inverse routine is

```
1   template <typename Type>
2   void apply_inverse_permutation(ulong (*x)(ulong), const Type *f, Type * restrict g, ulong n)
3   // Set g[x(k)] <-- f[k]  \forall k
4   // Must have:  0<=x(k)<n \forall k
5   {
6       for (ulong k=0; k<n; ++k)  g[x(k)] = f[k];
7   }
```

The in-place versions of these routines are identical to the routines that apply permutations given as arrays. Only a tiny change must be made in the processing of the cycles. For example, the fragment

```
1   void apply_permutation(const ulong *x, Type * restrict f, ulong n, bitarray *bp=0)
```

```
    [--snip--]
          ulong i = k;  // start of cycle
          Type t = f[i];
          ulong g = x[i];  // <--=
          while ( 0==(tp->test_set(g)) )
          {
              f[i] = f[g];
              i = g;
              g = x[i];  // <--=
          }
          f[i] = t;
    [--snip--]
```

must be changed to (replace '`x[i]`' by '`x(i)`')

```
  void apply_permutation(ulong (*x)(ulong), Type *f, ulong n, bitarray *bp=0)
    [--snip--]
          ulong i = k;  // start of cycle
          Type t = f[i];
          ulong g = x(i);  // <--=
          while ( 0==(tp->test_set(g)) )
          {
              f[i] = f[g];
              i = g;
              g = x(i);  // <--=
          }
          f[i] = t;
    [--snip--]
```

2.11.9 Random permutations

Routines for random permutations are given in [FXT: perm/permrand.h]. The following routine randomly permutes an array with arbitrary elements:

```
template <typename Type>
void random_permute(Type *f, ulong n)
{
    for (ulong k=1; k<n; ++k)
    {
        ulong r = (ulong)rand();
        r ^= r>>16;  // avoid using low bits of rand alone

        ulong i = r % (k+1);
        swap2(f[k], f[i]);
    }
}
```

The method is given in [108]. A random permutation can be obtained by applying the function to the canonical sequence:

```
void random_permutation(ulong *f, ulong n)
// Create a random permutation
{
    for (ulong k=0; k<n; ++k)  f[k] = k;
    random_permute(f, n);
}
```

We note that a slight modification of the underlying idea can be used for a routine for random selection from a list with only one linear read. Let L be a list of n items $L_1, \ldots, L_n$.

1. Set $t = L_1$, set $k = 1$.
2. Set $k = k + 1$. If $k > n$ return t.
3. With probability $1/k$ set $t = L_k$.
4. Go to step 2.

Note that one does not need to know n, the number of elements in the list, in advance: replace part 2 of step 2 by "If there are no more elements then return t".

A routine to apply a random cyclic permutation (as defined in section 2.11.4 on page 116) to an array is

```
template <typename Type>
void random_permute_cyclic(Type *f, ulong n)
// Permute the elements of f by a random cyclic permutation.
{
    for (ulong k=n-1; k>0; --k)
    {
        ulong r = (ulong)rand();
        r ^= r>>16;  // avoid using low bits of rand alone
        ulong i = r % k;
        swap2(f[k], f[i]);
    }
}
```

The method is described in [137], see also [277]. A random cyclic permutation can be obtained by applying a random cyclic permutation to the canonical sequence:

```
inline void
random_cyclic_permutation(ulong *f, ulong n)
// Create a random permutation that is cyclic.
{
    for (ulong k=0; k<n; ++k)  f[k] = k;
    random_permute_cyclic(f, n);
}
```

The cycle representation of a cyclic permutation can be obtained by applying a random permutation to all elements (of the identical permutation) except for the first element.

Chapter 3

Sorting and searching

In this chapter some practical flavors of sorting algorithms are given. These include plain sorting, sorting index arrays, pointer sorting; all optionally with a supplied comparison function. Massive literature exist about the topic so we will not go into the algorithmic details. Very readable texts are [95] and [236], in-depth information can be found in [169]. The sorting algorithms used in this chapter are selection sort, quicksort, counting sort and radix sort.

Some algorithms on sorted arrays like binary searching and determination of unique elements are included. Finally, some functions for scanning unsorted arrays are given.

3.1 Sorting

Selection sort

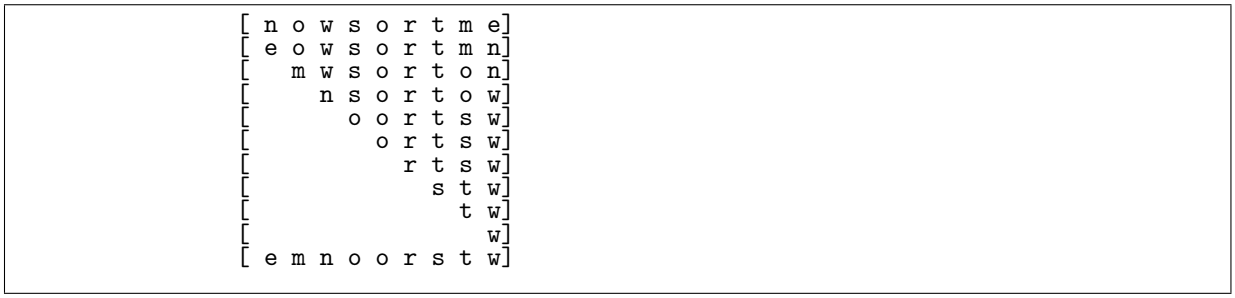

Figure 3.1-A: Sorting the string '`nowsortme`' with the selection sort algorithm.

There are a several algorithms for sorting that scale with $\sim n^2$ where n is the size of the array to be sorted. Here we use *selection sort* whose idea is to find the minimum of the array, swap it with the first element and repeat for all elements but the first. A demonstration of the algorithm is shown in figure 3.1-A, this is the output of [FXT: sort/selection-sort-demo.cc]. The implementation is straightforward [FXT: sort/sort.h]:

```
1   template <typename Type>
2   void selection_sort(Type *f, ulong n)
3   // Sort f[] (ascending order).
4   // Algorithm is proportional to O(n*n), use for short arrays only.
5   {
6       for (ulong i=0; i<n; ++i)
7       {
8           Type v = f[i];
9           ulong m = i; // position of minimum
10          ulong j = n;
```

```
11          while ( --j > i )  // search (index of) minimum
12          {
13              if ( f[j]<v )
14              {
15                  m = j;
16                  v = f[m];
17              }
18          }
19
20          swap2(f[i], f[m]);
21      }
22  }
```

A verification routine is always handy:

```
1   template  <typename Type>
2   bool is_sorted(const Type *f, ulong n)
3   // Return whether the sequence f[0], f[1], ..., f[n-1]
4   // is sorted in ascending order.
5   {
6       for (ulong k=1; k<n; ++k)  if ( f[k-1] > f[k] )  return false;
7       return  true;
8   }
```

A test for descending order is

```
1   template  <typename Type>
2   bool is_falling(const Type *f, ulong n)
3   {
4       for (ulong k=1; k<n; ++k)  if ( f[k-1] < f[k] )  return false;
5       return  true;
6   }
```

Quicksort

The *quicksort* algorithm scales $\sim n \log(n)$ (in the average case). It does not just obsolete the more simple schemes because for arrays small enough the 'simple' algorithm is usually the fastest method because of its minimal bookkeeping overhead, and it can be used inside the quicksort for lengths below some threshold.

The main ingredient of quicksort is to *partition* the array. The corresponding routine reorders the array and returns an *pivot* index p so that $\max(f_0, \ldots, f_{p-1}) \leq \min(f_p, \ldots, f_{n-1})$ [FXT: sort/sort.h]:

```
1   template <typename Type>
2   ulong partition(Type *f, ulong n)
3   {
4       // Avoid worst case with already sorted input:
5       const Type v = median3(f[0], f[n/2], f[n-1]);
6
7       ulong i = 0UL - 1;
8       ulong j = n;
9       while ( 1 )
10      {
11          do  { ++i; }  while ( f[i]<v );
12          do  { --j; }  while ( f[j]>v );
13
14          if ( i<j )  swap2(f[i], f[j]);
15          else        return j;
16      }
17  }
```

The function `median3()` is defined in [FXT: sort/minmaxmed23.h]:

```
1   template <typename Type>
2   static inline Type median3(const Type &x, const Type &y, const Type &z)
3   // Return median of the input values
4   { return  x<y ? (y<z ? y : (x<z ? z : x)) : (z<y ? y : (z<x ? z : x)); }
```

The function does 2 or 3 comparisons, depending on the input. One could simply use the element `f[0]` as pivot. However, the algorithm will be $\sim n^2$ (that is, quadratic) when the array is already sorted.

Quicksort calls `partition` on the whole array, then on the two parts left and right from the partition index, then for the four, eight, ... parts, until the parts are of length one. Note that the sub-arrays are usually of different lengths.

```
template <typename Type>
void quick_sort(Type *f, ulong n)
{
    if ( n<=1 )  return;

    ulong p = partition(f, n);
    ulong ln = p + 1;
    ulong rn = n - ln;
    quick_sort(f, ln);  // f[0]  ... f[ln-1]  left
    quick_sort(f+ln, rn);  // f[ln] ... f[n-1]   right
}
```

The actual implementation uses two optimizations: Firstly, when the size of the subproblems is smaller than a certain threshold selection sort is used. Secondly, the recursive calls are made for the smaller of the two sub-arrays, thereby the stack size is bounded by $\lceil \log_2(n) \rceil$.

```
template <typename Type>
void quick_sort(Type *f, ulong n)
{
 start:
    if ( n<8 ) // parameter: threshold for nonrecursive algorithm
    {
        selection_sort(f, n);
        return;
    }

    ulong p = partition(f, n);
    ulong ln = p + 1;
    ulong rn = n - ln;

    if ( ln>rn )  // recursion for shorter sub-array
    {
        quick_sort(f+ln, rn);  // f[ln] ... f[n-1]   right
        n = ln;
    }
    else
    {
        quick_sort(f, ln);  // f[0]  ... f[ln-1]  left
        n = rn;
        f += ln;
    }

    goto start;
}
```

The quicksort algorithm *will* be quadratic with certain inputs. A clever method to construct such inputs is described in [197]. A *heapsort* algorithm is $\sim n \cdot \log(n)$ also in the worst case, it is described in section 3.10 on page 142. Inputs that lead to quadratic time for the quicksort algorithm with median-of-3 partitioning are described in [206]. There it is suggested to use quicksort but detect problematic behavior during runtime and switch to heapsort if needed. The corresponding algorithm is called *introsort* (for *introspective Sorting*).

3.2 Binary search

The main reason for sorting may be that a fast search has to be performed repeatedly. The *binary search* algorithm works by the obvious subdivision of the data [FXT: `bsearch()` in sort/bsearch.h]:

```
template <typename Type>
ulong bsearch(const Type *f, ulong n, const Type v)
// Return index of first element in f[] that equals v
// Return ~0 if there is no such element.
// f[] must be sorted in ascending order.
// Must have  n!=0
{
    ulong nlo=0, nhi=n-1;
    while ( nlo != nhi )
```

```
10      {
11          ulong t = (nhi+nlo)/2;
12
13          if ( f[t] < v )  nlo = t + 1;
14          else             nhi = t;
15      }
16
17      if ( f[nhi]==v )  return nhi;
18      else              return ~0UL;
19  }
```

The algorithm uses $\sim \log_2(n)$ operations. For very large arrays the algorithm can be improved by selecting the new index `t` different from midpoint `(nhi+nlo)/2`, dependent of the value sought and the distribution of the values in the array. As a simple example consider an array of floating point numbers that are equally distributed in the interval $[\min(v), \max(v)]$. If the sought value equals v one would want to use the relation

$$\frac{n - \min(n)}{\max(n) - \min(n)} \quad = \quad \frac{v - \min(v)}{\max(v) - \min(v)} \tag{3.2-1}$$

where n denotes an index, and $\min(n), \max(n)$ denote the minimal and maximal index of the current interval. Solving for n gives the linear interpolation formula

$$n \quad = \quad \min(n) + \frac{\max(n) - \min(n)}{\max(v) - \min(v)} \, (v - \min(v)) \tag{3.2-2}$$

The corresponding *interpolation binary search* algorithm would select the new subdivision index `t` according to the given relation. One could even use quadratic interpolation schemes for the selection of `t`. For the majority of practical applications the midpoint version of the binary search will be good enough.

A simple modification of `bsearch` makes it search the first element greater than or equal to `v`: replace the operator `==` in the above code by `>=` and you have it: [FXT: `bsearch_ge()` in sort/bsearch.h]. Similar for the '`<=`' relation: `bsearch_le()`.

Approximate matches are found by [FXT: `bsearch_approx()` in sort/bsearchapprox.h]:

```
1   template <typename Type>
2   ulong bsearch_approx(const Type *f, ulong n, const Type v, Type da)
3   // Return index of first element x in f[] for which  |(x-v)| <= da
4   // Return ~0 if there is no such element.
5   // f[] must be sorted in ascending order.
6   // da must be positive.
7   //
8   // Makes sense only with inexact types (float or double).
9   // Must have  n!=0
10  {
11      ulong k = bsearch_ge(f, n, v-da);
12      if ( k<n )  k = bsearch_le(f+k, n-k, v+da);
13      return k;
14  }
```

3.3 Index sorting

While the 'plain' sorting reorders an array f so that, after it has finished, $f_k \le f_{k+1}$ the following routines sort an array of indices without modifying the actual data. The index-sort routines reorder the indices in an array `x` such that `x` applied to `f` as a permutation (in the sense of section 2.11.8 on page 119) will render `f` a sorted array [FXT: sort/sortidx.h]:

```
1   template <typename Type>
2   void idx_selection_sort(const Type *f, ulong n, ulong *x)
3   // Sort x[] so that the sequence
4   // f[x[0]], f[x[1]], ... f[x[n-1]]
5   // is sorted in ascending order.
6   // Algorithm is proportional to O(n*n), use for short array only.
7   {
```

```
8       for (ulong i=0; i<n; ++i)
9       {
10          Type v = f[x[i]];
11          ulong m = i; // position-ptr of minimum
12          ulong j = n;
13          while ( --j > i )  // search (index of) minimum
14          {
15              if ( f[x[j]]<v )
16              {
17                  m = j;
18                  v = f[x[m]];
19              }
20          }
21
22          swap2(x[i], x[m]);
23      }
24  }
```

Apart from the 'read only'-feature the index-sort routines have the nice property to perfectly work on non-contiguous data. The verification code is

```
1   template  <typename Type>
2   bool is_idx_sorted(const Type *f, ulong n, const ulong *x)
3   // Return whether the sequence
4   // f[x[0]], f[x[1]], ... f[x[n-1]]
5   // is sorted in ascending order.
6   {
7       if ( 0==n )  return 1;
8
9       while ( --n )  // n-1 ... 1
10      {
11          if ( f[x[n]] < f[x[n-1]] )  break;
12      }
13
14      return  !n;
15  }
```

The transformation of the `partition()` routine is straightforward:

```
1   template <typename Type>
2   ulong idx_partition(const Type *f, ulong n, ulong *x)
3   // rearrange index array, so that for some index p
4   // max(f[x[0]] ... f[x[p]]) <= min(f[x[p+1]] ... f[x[n-1]])
5   {
6       // Avoid worst case with already sorted input:
7       const Type v = median3(*x[0], *x[n/2], *x[n-1], cmp);
8
9       ulong i = 0UL - 1;
10      ulong j = n;
11      while ( 1 )
12      {
13          do  ++i;
14          while ( f[x[i]]<v );
15
16          do  --j;
17          while ( f[x[j]]>v );
18
19          if ( i<j )  swap2(x[i], x[j]);
20          else        return j;
21      }
22  }
```

The index-quicksort itself deserves a minute of contemplation comparing it to the plain version:

```
1   template <typename Type>
2   void idx_quick_sort(const Type *f, ulong n, ulong *x)
3   // Sort x[] so that the sequence
4   // f[x[0]], f[x[1]], ... f[x[n-1]]
5   // is sorted in ascending order.
6   {
7    start:
8       if ( n<8 ) // parameter: threshold for nonrecursive algorithm
9       {
10          idx_selection_sort(f, n, x);
11          return;
12      }
13
14      ulong p = idx_partition(f, n, x);
```

```
15      ulong ln = p + 1;
16      ulong rn = n - ln;
17
18      if ( ln>rn )  // recursion for shorter sub-array
19      {
20          idx_quick_sort(f, rn, x+ln);  // f[x[ln]] ... f[x[n-1]]  right
21          n = ln;
22      }
23      else
24      {
25          idx_quick_sort(f, ln, x);  // f[x[0]]  ... f[x[ln-1]]  left
26
27          n = rn;
28          x += ln;
29      }
30
31      goto start;
32  }
```

The index-analogues of the binary search algorithms are again straightforward, they are given in [FXT: sort/bsearchidx.h].

3.4 Pointer sorting

Pointer sorting is an idea similar to index sorting which is even less restricted than index sort: The data may be unaligned in memory. And overlapping. Or no data at all but port addresses controlling some highly dangerous machinery. Thereby pointer sort is the perfect way to highly cryptic and powerful programs that seg-fault when you least expect it.

Just to make the idea clear, the array of indices is replaced by an array of pointers [FXT: sort/sortptr.h].

```
1   template <typename Type>
2   void ptr_selection_sort(/*const Type *f,*/ ulong n, const Type **x)
3   // Sort x[] so that the sequence
4   // *x[0], *x[1], ..., *x[n-1]
5   // is sorted in ascending order.
6   // Algorithm is proportional to O(n*n), use for short arrays only.
7   {
8       for (ulong i=0; i<n; ++i)
9       {
10          Type v = *x[i];
11          ulong m = i; // position-ptr of minimum
12          ulong j = n;
13          while ( --j > i )  // search (index of) minimum
14          {
15              if ( *x[j]<v )
16              {
17                  m = j;
18                  v = *x[m];
19              }
20          }
21          swap2(x[i], x[m]);
22      }
23  }
```

The first argument (`const Type *f`) is not necessary with pointer sorting, it is indicated as comment to make the argument structure uniform. The verification routine is

```
1   template  <typename Type>
2   bool is_ptr_sorted(/*const Type *f,*/ ulong n, Type const*const*x)
3   // Return whether the sequence
4   // *x[0], *x[1], ..., *x[n-1]
5   // is sorted in ascending order.
6   {
7       if ( 0==n )  return 1;
8
9       while ( --n )  // n-1 ... 1
10      {
11          if ( *x[n] < *x[n-1] )  break;
12      }
13
```

```
14      return  !n;
15  }
```

The pointer versions of the search routines are given in [FXT: sort/bsearchptr.h].

3.5 Sorting by a supplied comparison function

The routines in [FXT: sort/sortfunc.h] are similar to the C-quicksort `qsort` that is part of the standard library. A comparison function `cmp` has to be supplied by the called so that compound data types can be sorted with respect to some key contained. Citing the manual page for `qsort`:

> The comparison function must return an integer less than, equal to, or greater than zero if the first argument is considered to be respectively less than, equal to, or greater than the second. If two members compare as equal, their order in the sorted array is undefined.

Note that the numerous calls to `cmp` do have a negative impact on the performance. With C++ you can provide a comparison 'function' for compound data by overloading the operators `<`, `<`, `<=` and `>=` and use the plain version. That is, the comparisons are inlined an we are back in performance land. Isn't C++ nice? As a prototypical example we give the selection sort routine:

```
1   template <typename Type>
2   void selection_sort(Type *f, ulong n, int (*cmp)(const Type &, const Type &))
3   // Sort f[] (ascending order)
4   // with respect to comparison function cmp().
5   // Algorithm is proportional to O(n*n), use for short array only.
6   {
7       for (ulong i=0; i<n; ++i)
8       {
9           Type v = f[i];
10          ulong m = i; // position of minimum
11          ulong j = n;
12          while ( --j > i )  // search (index of) minimum
13          {
14              if ( cmp(f[j],v) < 0 )
15              {
16                  m = j;
17                  v = f[m];
18              }
19          }
20
21          swap2(f[i], f[m]);
22      }
23  }
```

The other routines are rather straightforward translations of the (plain-) sort analogues: replace the comparison operations as follows

```
(a < b)      cmp(a,b) < 0
(a > b)      cmp(a,b) > 0
(a == b)     cmp(a,b) == 0
(a <= b)     cmp(a,b) <= 0
(a >= b)     cmp(a,b) >= 0
```

For example, the verification routine is

```
1   template  <typename Type>
2   bool is_sorted(const Type *f, ulong n, int (*cmp)(const Type &, const Type &))
3   // Return whether the sequence
4   // f[0], f[1], ..., f[n-1]
5   // is sorted in ascending order
6   // with respect to comparison function cmp().
7   {
8       if ( 0==n )  return 1;
9
10      while ( --n )  // n-1 ... 1
11      {
12          if ( cmp(f[n], f[n-1]) < 0 )  break;
13      }
14
15      return  !n;
```

```
16  }
```

3.5.1 Sorting complex numbers

You want to sort complex numbers? Fine for me, but *don't* tell your local mathematician. To see the mathematical problem we ask whether i is smaller or greater than zero. Assume $i > 0$: follows $i \cdot i > 0$ (we multiplied with a positive value) which is $-1 > 0$ and that is false. So, is $i < 0$? Then $i \cdot i > 0$ (multiplication with a negative value, as assumed). So $-1 > 0$, oops! The lesson is that there is no way to impose an arrangement on the complex numbers that would justify the usage of the symbols '$<$' and '$>$' consistent with the rules to manipulate inequalities.

Nevertheless we can invent a relation that allows us to sort: arranging (sorting) the complex numbers according to their absolute value (modulus) leaves infinitely many numbers in one 'bucket', namely all those that have the same distance from zero. However, one could use the modulus as the *major* ordering parameter, the angle as the *minor*. Or the real part as the major and the imaginary part as the minor. The latter is realized in

```
1   static inline int
2   cmp_complex(const Complex &f, const Complex &g)
3   {
4       const double fr = f.real(),  gr = g.real();
5       if ( fr!=gr )  return  (fr>gr ? +1 : -1);
6
7       const double fi = f.imag(),  gi = g.imag();
8       if ( fi!=gi )  return  (fi>gi ? +1 : -1);
9
10      return  0;
11  }
```

This routine, when used as comparison with the function-sort, as in

```
1   void complex_sort(Complex *f, ulong n)
2   // major order wrt. real part
3   // minor order wrt. imag part
4   {
5       quick_sort(f, n, cmp_complex);
6   }
```

can indeed be the practical tool you had in mind.

3.5.2 Index and pointer sorting

The index sorting routines that use a supplied comparison function are given in [FXT: sort/sortidxfunc.h]:

```
1   template <typename Type>
2   void idx_selection_sort(const Type *f, ulong n, ulong *x,
3                           int (*cmp)(const Type &, const Type &))
4   // Sort x[] so that the sequence
5   // f[x[0]], f[x[1]], ... f[x[n-1]]
6   // is sorted in ascending order
7   // with respect to comparison function cmp()
8   // Algorithm is proportional to O(n*n), use for short array only.
9   {
10      for (ulong i=0; i<n; ++i)
11      {
12          Type v = f[x[i]];
13          ulong m = i; // position-ptr of minimum
14          ulong j = n;
15          while ( --j > i )  // search (index of) minimum
16          {
17              if ( cmp(f[x[j]], v) < 0 )
18              {
19                  m = j;
20                  v = f[x[m]];
21              }
22          }
23
24          swap2(x[i], x[m]);
```

```
25      }
26  }
```

The verification routine is:

```
1   template  <typename Type>
2   bool is_idx_sorted(const Type *f, ulong n, const ulong *x,
3                      int (*cmp)(const Type &, const Type &))
4   // Return whether the sequence
5   // f[x[0]], f[x[1]], ... f[x[n-1]]
6   // is sorted in ascending order
7   // with respect to comparison function cmp()
8   {
9       if ( 0==n )  return 1;
10
11      while ( --n )  // n-1 ... 1
12      {
13          if ( cmp(f[x[n]], f[x[n-1]]) < 0 )  break;
14      }
15
16      return  !n;
17  }
```

The pointer sorting versions are given in [FXT: sort/sortptrfunc.h]

```
1   template <typename Type>
2   void ptr_selection_sort(/*const Type *f,*/ ulong n, const Type **x,
3                           int (*cmp)(const Type &, const Type &))
4   // Sort x[] so that the sequence
5   // *x[0], *x[1], ..., *x[n-1]
6   // is sorted in ascending order
7   // with respect to comparison function cmp().
8   // Algorithm is proportional to O(n*n), use for short array only.
9   {
10      for (ulong i=0; i<n; ++i)
11      {
12          Type v = *x[i];
13          ulong m = i; // position-ptr of minimum
14          ulong j = n;
15          while ( --j > i )  // search (index of) minimum
16          {
17              if ( cmp(*x[j],v)<0 )
18              {
19                  m = j;
20                  v = *x[m];
21              }
22          }
23
24          swap2(x[i], x[m]);
25      }
26  }
```

The verification routine is:

```
1   template  <typename Type>
2   bool is_ptr_sorted(/*const Type *f,*/ ulong n, Type const*const*x,
3                      int (*cmp)(const Type &, const Type &))
4   // Return whether the sequence
5   // *x[0], *x[1], ..., *x[n-1]
6   // is sorted in ascending order
7   // with respect to comparison function cmp().
8   {
9       if ( 0==n )  return 1;
10
11      while ( --n )  // n-1 ... 1
12      {
13          if ( cmp(*x[n],*x[n-1])<0 )  break;
14      }
15
16      return  !n;
17  }
```

The corresponding versions of the binary search algorithm are given in [FXT: sort/bsearchidxfunc.h] and [FXT: sort/bsearchptrfunc.h].

3.6 Determination of unique elements

We present functions that check whether values in a sorted array are repeated or unique. All routines are taken from [FXT: sort/unique.h]. To test whether all values are unique, use

```
template  <typename Type>
ulong test_unique(const Type *f, ulong n)
// For a sorted array test whether all values are unique
//   (i.e. whether no value is repeated).
// Return 0 if all values are unique else return index of the second
// element in the first pair found.
{
    for (ulong k=1; k<n; ++k)
    {
        if ( f[k] == f[k-1] )  return  k;  // k != 0
    }

    return  0;
}
```

```
template  <typename Type>
ulong is_unique(const Type *f, ulong n)
// Return true if all values are unique, else return false.
{
    return ( 0==test_unique(f, n) );
}
```

Counting the elements that appear just once:

```
template  <typename Type>
int unique_count(const Type *f, ulong n)
// For a sorted array return the number of unique values
// the number of (not necessarily distinct) repeated
//   values is n - unique_count(f, n);
{
    if ( 1>=n )  return n;

    ulong ct = 1;
    for (ulong k=1; k<n; ++k)
    {
        if ( f[k] != f[k-1] )  ++ct;
    }

    return  ct;
}
```

Removing repeated elements:

```
template  <typename Type>
ulong unique(Type *f, ulong n)
// For a sorted array squeeze all repeated values
// and return the number of unique values.
// Example:  [1, 3, 3, 4, 5, 8, 8] --> [1, 3, 4, 5, 8]
// The routine also works for unsorted arrays as long
// as identical elements only appear in contiguous blocks.
// Example: [4, 4, 3, 7, 7] --> [4, 3, 7]
// The order is preserved.
{
    ulong u = unique_count(f, n);

    if ( u==n )  return n;  // nothing to do

    Type v = f[0];
    for (ulong j=1, k=1;  j<u;  ++j)
    {
        while ( f[k]==v )  ++k;  // search next different element
        v = f[j] = f[k];
    }

    return u;
}
```

The inner `while`-loop does never access an element out of bounds as it is executed only as long as there is at least one remaining change of value inside the array.

3.7 Unique elements with inexact types

Determination of unique elements with inexact types (floats) is a bit tricky as one cannot rely that elements that *should* be identical are exactly equal. A solution to the problem is to allow for a maximal (absolute) difference within which two contiguous elements will still be considered equal can be provided as additional parameter. We replace equality conditions with a call to [FXT: sort/uniqueapprox.h]

```
template  <typename Type>
inline bool approx_equal(Type x1, Type x2, Type da)
// Return whether abs(x2-x1) <= da
// Must have da>=0
{
    Type d = x2 - x1;
    if ( d<=0 )  d = -d;
    if ( d <= da )  return true;
    else            return false;
}
```

The verification routine is

```
template  <typename Type>
ulong test_unique_approx(const Type *f, ulong n, Type da)
// For a sorted array test whether all values are
// unique within some tolerance (i.e. whether no value is repeated).
// Return 0 if all values are unique,
// else return index of the second element in the first pair found.
// Makes mostly sense with inexact types (float or double)
{
    if ( da<=0 )  da = -da;  // want positive tolerance

    for (ulong k=1; k<n; ++k)
    {
        if ( approx_equal(f[k], f[k-1], da) )  return k;  // k != 0
    }

    return  0;
}
```

One subtle point is that the values can slowly 'drift away' unnoticed by this implementation: consider a long array where each difference computed has the same sign and is just smaller than `da`, say it is $d = 0.6 \cdot$`da`. The difference of the first and last value then is $0.6 \cdot (n-1) \cdot d$ which is greater than `da` for $n \geq 3$.

The number of unique elements can be counted as follows:

```
template  <typename Type>
ulong unique_approx_count(const Type *f, ulong n, Type da)
// For a sorted array return the number of unique values
// the number of (not necessarily distinct) repeated
//   values is n - unique_approx_count(f, n, da);
{
    if ( 1>=n )  return n;
    if ( da<=0 )  da = -da;  // Must have positive tolerance
    ulong ct = 1;
    for (ulong k=1; k<n; ++k)
    {
        if ( approx_equal(f[k], f[k-1], da) )  ++ct;
    }
    return  ct;
}
```

The following routine removes duplicates:

```
template  <typename Type>
ulong unique_approx(Type *f, ulong n, Type da)
// For a sorted array squeeze all repeated (within tolerance da) values
// and return the number of unique values.
// Example:  [1, 3, 3, 4, 5, 8, 8] --> [1, 3, 4, 5, 8]
// The routine also works for unsorted arrays as long
// as identical elements only appear in contiguous blocks.
// Example: [4, 4, 3, 7, 7] --> [4, 3, 7]
// The order is preserved.
{
```

```
11      ulong u = unique_approx_count(f, n, da);
12
13      if ( u==n )  return n;  // nothing to do
14
15      if ( da<=0 )  da = -da;  // Must have positive tolerance
16      Type v = f[0];
17      for (ulong j=1, k=1;  j<u;  ++j)
18      {
19          // search next different element:
20          while ( approx_equal(f[k], v, da) )
21          {
22              v = f[k];  // avoid problem with slowly drifting values
23              ++k;
24          }
25          v = f[j] = f[k];
26      }
27
28      return u;
29  }
```

A useful preprocessing step (before using `test_unique_approx()`) is to quantize the elements of an array [FXT: `quantize()` in sort/quantize.h]:

```
1   template  <typename Type>
2   void quantize(Type *f, ulong n, double q)
3   // In f[] set each element x to q*floor(1/q*(x+q/2))
4   // E.g.: q=1 ==> round to nearest integer
5   //       q=1/1000 ==> round to nearest multiple of 1/1000
6   // For inexact types (float or double).
7   {
8       Type qh = q * 0.5;
9       Type q1 = 1.0 / q;
10      while ( n-- )
11      {
12          f[n] = q * floor( q1 * (f[n]+qh) );
13      }
14  }
```

One should use a quantization parameter `q` that is greater than the value used for `da`.

A simple demonstration is given in [FXT: sort/unique-demo.cc]:

```
Random values:
   0:  0.9727750243
   1:  0.2925167845
   2:  0.7713576982
   3:  0.5267449795
   4:  0.7699138366
   5:  0.4002286223

Quantization with q=0.01

Quantized & sorted :
   0:  0.2900000000
   1:  0.4000000000
   2:  0.5300000000
   3:  0.7700000000
   4:  0.7700000000
   5:  0.9700000000
First REPEATED value at index 4  (and 3)

Unique'd array:
   0:  0.2900000000
   1:  0.4000000000
   2:  0.5300000000
   3:  0.7700000000
   4:  0.9700000000
```

The routine `quantize()` turns out to be also useful for the conversion of imprecise data to symbols. For example, the array of floating point values on the left corresponds to the symbolic (numbers used as symbols) table on the right:

```
 1.3133     -1.0101     0.79412     -0.71544          9   2   6   3
 0.29064     0.99173   -1.4382       0.79412          5   7   0   6
-1.1086      1.2521     0.99173     -1.0101           1   8   7   2
-0.18003    -1.1086     0.29064      1.3133           4   1   5   9
```

In this example values were considered identical when their absolute difference is less than 10^{-3}. The symbolic representation can be helpful to recognize structure in imprecise data. The routine is [FXT: sort/symbolify.h]:

```
1   template <typename Type>
```

```
2   ulong symbolify_by_size(const Type *f, Type * restrict g, ulong n,
3                           Type eps=1e-6, ulong *ix=0)
4   // From f[] compute an array of 'symbols' g[] (i.e. numbers)
5   //   that represent the different values.
6   // Values are considered identical if their absolute difference
7   //   is less than eps.
8   // Symbols are given with respect to sort-order.
9   // Return number of different values found (after quantize).
10  // Optionally supply x[] (scratch space for permutations).
11  {
12      ulong *x = ix;
13      if ( 0==ix )  x = new ulong[n];
14
15      set_seq(x, n);
16      idx_quick_sort(f, n, x);
17
18      apply_permutation(x, f, g, n);
19      quantize(g, n, eps);
20
21      eps *= 0.5;  // some val <1.0
22      ulong nsym = 1;
23
24      ulong z = 0;
25      Type s = 0.0;
26      Type el = g[z], lel;
27      g[z] = s;
28      for(ulong k=z+1; k<n; ++k)
29      {
30          lel = el;
31          el = g[k];
32          if ( fabs(el-lel) > eps )
33          {
34              ++nsym;
35              s += 1.0;
36          }
37          g[k] = s;
38      }
39
40      apply_inverse_permutation(x, g, n);
41
42      if ( 0==ix )  delete [] x;
43
44      return  nsym;
45  }
```

The example shown was created with the program [FXT: sort/symbolify-demo.cc]. The routines `apply_permutation()` and `apply_inverse_permutation()` are given in section 2.11.8 on page 119.

3.8 Determination of equivalence classes

Let S be a set and $C := S \times S$ the set of all ordered pairs (x, y) with $x, y \in S$. A *binary relation* R on S is a subset of C. An *equivalence relation* is a binary relation that has three additional properties:

- *reflexive*: $x \equiv x \; \forall x$.
- *symmetric*: $x \equiv y \iff y \equiv x$.
- *transitive*: $x \equiv y,\ y \equiv z \implies x \equiv z$.

Here we wrote $x \equiv y$ for $(x, y) \in R$ where $x, y \in S$.

We want to determine the *equivalence classes*: an equivalence relation partitions a set into $1 \le q \le n$ subsets E_1, E_2, ..., E_q so that $x \equiv y$ whenever both x and y are in the same subset but $x \not\equiv y$ if x and y are in different subsets.

For example, the usual equality relation is an equivalence relation, with a set of (different) numbers each number is in its own class. With the equivalence relation that $x \equiv y$ whenever $x - y$ is a multiple of

some fixed integer m and the set $\mathbb{Z}$ of all natural numbers we obtain m subsets and $x \equiv y$ if and only if $x \equiv y \bmod m$.

Let n be the number of elements in S and Q be a set so that, on termination of the algorithm, $Q_k = j$ if j is the least index so that $S_j \equiv S_k$ (note that we consider the sets to be in a fixed but arbitrary order here).

We proceed as follows:

1. Put each element in its own equivalence class: $Q_k := k$ for all $0 \le k < n$
2. Set $k := 1$ (index of the second element).
3. Search downwards for an equivalent element: for $j = k-1, \ldots, 0$ test whether $S_k \equiv S_j$. If so (at position j), set $Q_k = Q_j$ and goto step 4.
4. Set $k := k+1$ and if $k < n$ goto step 3, else terminate.

We can terminate the search with the first equivalent element found because, if j is the index of the equivalent, the Q_j is already minimal.

The lower and upper bounds for the computational cost are n and n^2, respectively: the algorithm needs proportional n operations when all elements are in the same equivalence class and n^2 operations when each element lies in its own class.

A C++ implementation is [FXT: `equivalence_classes()` in sort/equivclasses.h]. The equivalence relation must be supplied as a function `equiv_q()` that returns true when its arguments are equivalent.

```
1   template <typename Type>
2   void equivalence_classes(const Type *s, ulong n, bool (*equiv_q)(Type,Type), ulong *q)
3   // Given an equivalence relation '==' (as function equiv_q())
4   //  and a set s[] with n elements,
5   // write into q[k] (0<=k<n) the index j of the
6   //  first element s[j] so that s[k]==s[j].
7   // For the complexity C:  n<=C<=n*n
8   //  C=n*n if each element is in its own class
9   //  C=n   if all elements are in the same class
10  {
11      for (ulong k=0; k<n; ++k)  q[k] = k;  // each in own class
12      for (ulong k=1; k<n; ++k)
13      {
14          ulong j = k;
15          while ( j-- )
16          {
17              if ( equiv_q(s[j], s[k]) )
18              {
19                  q[k] = q[j];
20                  break;
21              }
22          }
23      }
24  }
```

3.8.1 Examples for equivalence classes

3.8.1.1 Integers modulo m

Choose an integer $m \ge 2$ and let any two integers a and b be equivalent if $a - b$ is an integer multiple of m. We can choose the numbers $0, 1 \ldots, m-1$ as *representatives* of the m classes obtained. Now we can do computations with those classes via the modular arithmetic as described in section 37.1 on page 753. This is easily the most important example of all equivalence classes.

We note that the concept still make sense with a real (that is, possibly non-integral) modulus m. We still put two numbers a and b into the same class if $a - b$ is a *integer* multiple of m. Finally, the modulus zero leads to the equivalence relation 'equality'.

3.8.1.2 Binary necklaces

The set S of n-bit binary words and the equivalence relation so that two words x and y are equivalent when there is a cyclic shift $h_k(x)$ by $0 \le k < n$ positions so that $h_k(x) = y$. The relation is supplied as the function [FXT: sort/equivclass-necklaces-demo.cc]:

```
1   static ulong b;  // number of bits
2   bool n_equiv_q(ulong x, ulong y)  // necklaces
3   {
4       ulong d = bit_cyclic_dist(x, y, b);
5       return  (0==d);
6   }
```

With $n = 4$ we obtain the following list of equivalence classes:

```
     0:   ....  [#=1]
     1:   1...  .1..  ...1  ..1.  [#=4]
     3:   1..1  11..  ..11  .11.  [#=4]
     5:   .1.1  1.1.  [#=2]
     7:   11.1  111.  1.11  .111  [#=4]
    15:   1111  [#=1]
  # of equivalence classes = 6
```

These correspond to the *binary necklaces* of length 4. One usually chooses the cyclic minima (or maxima) among equivalent words as representatives of the classes.

3.8.1.3 Unlabeled binary necklaces

Same set but the equivalence relation is defined to identify two words x and y when there is a cyclic shift $h_k(x)$ by $0 \le k < b$ positions so that either $h_k(x) = y$ or $h_k(x) = \overline{y}$ where $\overline{y}$ is the complement of y:

```
1   static ulong mm;  // mask to complement
2   bool nu_equiv_q(ulong x, ulong y)  // unlabeled necklaces
3   {
4       ulong d = bit_cyclic_dist(x, y, b);
5       if ( 0!=d )  d = bit_cyclic_dist(mm^x, y, b);
6       return  (0==d);
7   }
```

With $n = 4$ we obtain:

```
    0:   1111  ....  [#=2]
    1:   111.  11.1  1.11  1...  .111  ...1  ..1.  .1..  [#=8]
    3:   .11.  1..1  11..  ..11  [#=4]
    5:   .1.1  1.1.  [#=2]
  # of equivalence classes = 4
```

These correspond to the *unlabeled binary necklaces* of length 4.

3.8.1.4 Binary bracelets

The *binary bracelets* are obtained by identifying two words that are identical up to rotation and possible reversion. The corresponding comparison function is

```
1   bool b_equiv_q(ulong x, ulong y)  // bracelets
2   {
3       ulong d = bit_cyclic_dist(x, y, b);
4       if ( 0!=d )  d = bit_cyclic_dist(revbin(x,b), y, b);
5       return  (0==d);
6   }
```

There are six binary bracelets of length 4:

```
  0:   ....  [#=1]
  1:   1...  .1..  ...1  ..1.  [#=4]
  3:   1..1  11..  ..11  .11.  [#=4]
  5:   .1.1  1.1.  [#=2]
  7:   11.1  111.  1.11  .111  [#=4]
 15:   1111  [#=1]
```

The *unlabeled binary bracelets* are obtained by additionally allowing for bit-wise complementation:

```
1   bool bu_equiv_q(ulong x, ulong y)  // unlabeled bracelets
2   {
3       ulong d = bit_cyclic_dist(x, y, b);
4       x ^= mm;
5       if ( 0!=d )  d = bit_cyclic_dist(x, y, b);
6
7       x = revbin(x,b);
8       if ( 0!=d )  d = bit_cyclic_dist(x, y, b);
9       x ^= mm;
10      if ( 0!=d )  d = bit_cyclic_dist(x, y, b);
11
12      return  (0==d);
13  }
```

There are four unlabeled binary bracelets of length 4:

```
0:   1111  ....  [#=2]
1:   111.  11.1  1.11  1...  .111  ...1  ..1.  .1..  [#=8]
3:   .11.  1..1  11..  ..11  [#=4]
5:   .1.1  1.1.  [#=2]
```

The shown functions are given in [FXT: sort/equivclass-bracelets-demo.cc] which can be used to produce listings of the equivalence classes.

We give the number of binary necklaces '`N`', bracelets '`B`', unlabeled necklaces '`N/U`' and unlabeled bracelets '`B/U`'. The second row gives the sequence number of [240].

n :	N	B	N/U	B/U
[240]#	A000031	A000029	A000013	A000011
1:	2	2	1	1
2:	3	3	2	2
3:	4	4	2	2
4:	6	6	4	4
5:	8	8	4	4
6:	14	13	8	8
7:	20	18	10	9
8:	36	30	20	18
9:	60	46	30	23
10:	108	78	56	44
11:	188	126	94	63
12:	352	224	180	122
13:	632	380	316	190
14:	1182	687	596	362
15:	2192	1224	1096	612

3.8.1.5 Binary words with reversion and complement

The set S of n-bit binary words and the equivalence relation identifying two words x and y whenever they are mutual complements or bit-wise reversals.

For example, the equivalence classes with 3-, 4- and 5-bit words are shown in figure 3.8-A. The sequence of numbers of equivalence classes for word-sizes n is (entry A005418 of [240])

```
n:  1, 2, 3, 4,  5,  6,  7,  8,   9,  10,  11,   12,   13,   14,   15,    16, ...
#:  1, 2, 3, 6, 10, 20, 36, 72, 136, 272, 528, 1056, 2080, 4160, 8256, 16512, ...
```

The equivalence classes can be computed with the program [FXT: sort/equivclass-bitstring-demo.cc].

We have chosen examples where the resulting equivalence classes can be verified by inspection. For example, we could create the subsets of equivalent necklaces by simply rotating a given word and marking the so far visited words. Such an approach, however, is not possible in general when the equivalence relation does not have an obvious structure.

```
      3 classes with 3-bit words:
        0:   111  ...
        1:   ..1  .11  1..  11.
        2:   1.1  .1.


      6 classes with 4-bit words:
        0:   1111  ....
        1:   111.  1...  .111  ...1
        2:   ..1.  .1..  1.11  11.1
        3:   11..  ..11
        5:   1.1.  .1.1
        6:   .11.  1..1
```

```
     10 classes with 5-bit words:
       0:   11111  .....
       1:   1111.  1....  .1111  ....1
       2:   1.111  111.1  .1...  ...1.
       3:   111..  ...11  ..111  11...
       4:   ..1..  11.11
       5:   11.1.  1.1..  ..1.1  .1.11
       6:   ..11.  .11..  11..1  1..11
       9:   .11.1  1.11.  .1..1  1..1.
      10:   .1.1.  1.1.1
      14:   1...1  .111.
```

Figure 3.8-A: Equivalence classes of binary words where words are identified if either their reversals or complements are equal.

3.8.2 The number of equivalence relations for a set of n elements

The sequence $B(n)$ of the number of possible partitionings (and thereby equivalence relations) for the set $\{1, 2, \ldots, n\}$ starts as ($n \geq 1$):

```
1, 2, 5, 15, 52, 203, 877, 4140, 21147, 115975, 678570, 4213597, ...
```

These are the *Bell numbers*, sequence A000110 of [240]. They can be computed easily as indicated in the following table:

```
0:  [ 1]
1:  [ 1,  2]
2:  [ 2,  3,  5]
3:  [ 5,  7, 10,  15]
4:  [15, 20, 27,  37,  52]
5:  [52, 67, 87, 114, 151, 203]
n:  [B(n), ... ]
```

The first element in each row is the last element of the previous row, the remaining elements are the sum of their left and upper left neighbors. As pari/gp code:

```
1  N=7;  v=w=b=vector(N);  v[1]=1;
2  { for(n=1,N-1,
3      b[n] = v[1];
4      print(n-1, ":  ", v); \\ print row
5      w[1] = v[n];
6      for(k=2,n+1, w[k]=w[k-1]+v[k-1]);
7      v=w;
8   ); }
```

An implementation in C++ is given in [FXT: comb/bell-number-demo.cc]. An alternative way to compute the Bell Numbers is shown in section 15.1 on page 336.

3.9 Determination of monotonicity and convexity *

A sequence is called *monotone* if it is either purely ascending or purely descending. This includes the case where subsequent elements are equal. Whether a constant sequence is considered ascending or descending in this context is a matter of convention.

A routine to check for monotonicity is [FXT: sort/monotone.h]:

```
1   template  <typename Type>
2   int is_monotone(const Type *f, ulong n)
3   // Return
4   //   +1 for ascending order
5   //   -1 for descending order
6   //   else 0
7   {
8       if ( 1>=n )  return  +1;
9
10      ulong k;
```

```
11      for (k=1; k<n; ++k)  // skip constant start
12      {
13          if ( f[k] != f[k-1] )  break;
14      }
15
16      if ( k==n )  return +1;  // constant is considered ascending here
17
18      int s = ( f[k] > f[k-1] ? +1 : -1 );
19
20      if ( s>0 ) // was: ascending
21      {
22          // scan for descending pair:
23          for (  ; k<n; ++k)  if ( f[k] < f[k-1] )  return 0;
24      }
25      else // was: descending
26      {
27          // scan for ascending pair:
28          for (  ; k<n; ++k)  if ( f[k] > f[k-1] )  return 0;
29      }
30
31      return  s;
32  }
```

A *strictly monotone* sequence is a monotone sequence that has no identical pairs of elements. The test turns out to be slightly easier:

```
1   template  <typename Type>
2   int is_strictly_monotone(const Type *f, ulong n)
3   // return
4   //   +1 for strictly ascending order
5   //   -1 for strictly descending order
6   //   else 0
7   {
8       if ( 1>=n )  return  +1;
9
10      ulong k = 1;
11      if ( f[k] == f[k-1] )  return 0;
12
13      int s = ( f[k] > f[k-1] ? +1 : -1 );
14
15      if ( s>0 ) // was: ascending
16      {
17          // scan for descending pair:
18          for (  ; k<n; ++k)  if ( f[k] <= f[k-1] )  return 0;
19      }
20      else // was: descending
21      {
22          // scan for ascending pair:
23          for (  ; k<n; ++k)  if ( f[k] >= f[k-1] )  return 0;
24      }
25
26      return  s;
27  }
```

A sequence is called *convex* if it starts with an ascending part and ends with a descending part. A *concave* sequence starts with a descending and ends with an ascending part. Whether a monotone sequence is considered convex or concave again is a matter of convention (you have the choice to consider the first or the last element as extremum). Lacking a term that contains both convex and concave the following routine is called `is_convex()` [FXT: sort/convex.h]:

```
1   template <typename Type>
2   long is_convex(Type *f, ulong n)
3   // Return
4   //   +val for convex sequence (first rising then falling)
5   //   -val for concave sequence (first falling then rising)
6   //   else 0
7   //
8   // val is the (second) index of the first pair at the point
9   // where the ordering changes;  val>=n iff seq. is monotone.
10  //
11  // Note: a constant sequence is considered any of rising/falling
12  {
13      if ( 1>=n )  return +1;
14
15      ulong k = 1;
16      for (k=1; k<n; ++k)  // skip constant start
17      {
```

```
18          if ( f[k] != f[k-1] )  break;
19      }
20
21      if ( k==n )  return  +n;  // constant is considered convex here
22
23      int s = ( f[k] > f[k-1] ? +1 : -1 );
24
25      if ( s>0 ) // was: ascending
26      {
27          // scan for strictly descending pair:
28          for (  ; k<n; ++k)  if ( f[k] < f[k-1] )  break;
29          s = +k;
30      }
31      else // was: descending
32      {
33          // scan for strictly ascending pair:
34          for (  ; k<n; ++k)  if ( f[k] > f[k-1] )  break;
35          s = -k;
36      }
37
38      if ( k==n )  return s;  // sequence is monotone
39
40      // check that the ordering does not change again:
41      if ( s>0 ) // was: ascending --> descending
42      {
43          // scan for strictly ascending pair:
44          for (  ; k<n; ++k)  if ( f[k] > f[k-1] )  return 0;
45      }
46      else // was: descending
47      {
48          // scan for strictly descending pair:
49          for (  ; k<n; ++k)  if ( f[k] < f[k-1] )  return 0;
50      }
51
52      return  s;
53  }
```

The test for *strictly convex* (or concave) sequences is:

```
1   template <typename Type>
2   long is_strictly_convex(Type *f, ulong n)
3   // Return
4   //   +val for strictly convex sequence
5   //        (i.e. first strictly rising then strictly falling)
6   //   -val for strictly concave sequence
7   //        (i.e. first strictly falling then strictly rising)
8   //   else 0
9   //
10  // val is the (second) index of the first pair at the point
11  // where the ordering changes;  val>=n iff seq. is strictly monotone.
12  {
13      if ( 1>=n )  return +1;
14
15      ulong k = 1;
16      if ( f[k] == f[k-1] )  return 0;
17
18      int s = ( f[k] > f[k-1] ? +1 : -1 );
19
20      if ( s>0 ) // was: ascending
21      {
22          // scan for descending pair:
23          for (  ; k<n; ++k)  if ( f[k] <= f[k-1] )  break;
24          s = +k;
25      }
26      else // was: descending
27      {
28          // scan for ascending pair:
29          for (  ; k<n; ++k)  if ( f[k] >= f[k-1] )  break;
30          s = -k;
31      }
32
33      if ( k==n )  return s;  // sequence is monotone
34      else  if ( f[k] == f[k-1] )  return 0;
35
36      // check that the ordering does not change again:
37      if ( s>0 ) // was: ascending --> descending
38      {
```

```
39          // scan for ascending pair:
40          for (  ; k<n; ++k)  if ( f[k] >= f[k-1] )  return 0;
41      }
42      else // was: descending
43      {
44          // scan for descending pair:
45          for (  ; k<n; ++k)  if ( f[k] <= f[k-1] )  return 0;
46      }
47
48      return  s;
49  }
```

3.10 Heapsort

The *heapsort* algorithm uses the heap data structure introduced in section 4.5 on page 157. A heap can be sorted by swapping the first (and biggest) element with the last and 'repairing' the array of size $n-1$ by a call to `heapify1()`. Applying this idea recursively until there is nothing more to sort leads to the routine [FXT: sort/heapsort.h]:

```
1   template <typename Type>
2   void heap_sort_ascending(Type *x, ulong n)
3   // Sort an array that has the heap-property into ascending order.
4   // On return x[] is _not_ a heap anymore.
5   {
6       Type *p = x - 1;
7       for (ulong k=n; k>1; --k)
8       {
9           swap2(p[1], p[k]);   // move largest to end of array
10          --n;                 // remaining array is one element less
11          heapify1(p, n, 1);  // restore heap-property
12      }
13  }
```

that needs time $O(n \log(n))$. That is, a call to

```
1   template <typename Type>
2   void heap_sort(Type *x, ulong n)
3   {
4       build_heap(x, n);
5       heap_sort_ascending(x, n);
6   }
```

will sort the array `x[]` into ascending order. Note that sorting into descending order is not any harder:

```
1   template <typename Type>
2   void heap_sort_descending(Type *x, ulong n)
3   // Sort an array that has the heap-property into descending order.
4   // On return x[] is _not_ a heap anymore.
5   {
6       Type *p = x - 1;
7       for (ulong k=n; k>1; --k)
8       {
9           ++p;  --n;           // remaining array is one element less
10          heapify1(p, n, 1);  // restore heap-property
11      }
12  }
```

A program that demonstrates the algorithm is [FXT: sort/heapsort-demo.cc].

3.11 Counting sort and radix sort

Imagine you want to sort an n-element array F of (unsigned) 8-bit values. An sorting algorithm that only uses 2 passes through the data proceeds as follows:

1. Allocate an array C of 256 integers and set all its elements to zero.

2. Count: for $k = 0, 1, \ldots, n-1$ increment $C[F[k]]$.
   Now $C[x]$ contains how many bytes in F have the value x.

3. Set $r = 0$. For $j = 0, 1, \ldots, 255$
   set $k = C[j]$ and write j to the elements $F[r]$, $F[r+1]$, $\ldots$, $F[r+k-1]$ then add k to r.

For large values of n this method is significantly faster than any other sorting algorithm. Note that no comparisons are made between the elements of F. Instead they are counted, the algorithm is the *counting sort* algorithm.

It might seem that the idea applies only to very special cases but with a little care it can be used in more general situations. We modify the method so that we are able to sort also (unsigned) integer variables whose range of values would make the method impractical with respect to a subrange of the bits in each word. We need an array G that has as many elements as F:

1. Choose any consecutive run of b bits, these will be represented by a bit mask m. Allocate an array C of 2^b integers and set all its elements to zero.

2. Let M be a function that maps the (2^b) values of interest (the bits masked out by m) to the range $0, 1, \ldots, 2^b - 1$.

3. Count: for $k = 0, 1, \ldots, n-1$ increment $C[M(F[k])]$.
   Now $C[x]$ contains how many values of $M(F[.])$ equal x.

4. Cumulate: for $j = 1, 2, \ldots, 2^b - 1$ (second to last) add $C[j-1]$ to $C[j]$.
   Now $C[x]$ contains the number of values $M(F[.])$ less or equal to x.

5. Copy: for $k = n-1, \ldots, 2, 1, 0$ (last to first) set $x := M(F[k])$, decrement $C[x]$ then set $i = C[x]$, then set $G[i] := F[x]$.

A crucial property of the algorithm is that it is *stable*: when we use it to sort with respect to a certain bitmask m and there is more than one element being mapped to the same value then the relative order between these elements is preserved.

```
Input                Counting sort wrt. two lowest bits
                        m = .....11
 0:   11111.11<       0:   ...i...
 1:   ...i...         1:   .iiii..
 2:   ..i.i.i         2:   .iii...
 3:   ..i....i        3:   ..i.i.i
 4:   .iiiii<         4:   .i..i..i
 5:   .iiii..         5:   ..i...i.
 6:   .i..i..i        6:   .i.i.ii.
 7:   .i.i.ii.        7:   11111.11<
 8:   .ii...ii<       8:   ..i.iiii<
 9:   .iii....        9:   .ii...11<
```

Note that relative order of the three words ending with two set bits (marked with '<') is preserved.

A routine that verifies whether an array is sorted with respect to a bit range specified by the variable `b0` and `m` is [FXT: sort/radixsort.cc]:

```
1   bool
2   is_counting_sorted(const ulong *f, ulong n, ulong b0, ulong m)
3   // Whether f[] is sorted wrt. bits b0,...,b0+z-1
4   //  where z is the number of bits set in m.
5   // m must contain a single run of bits starting at bit zero.
6   {
7       m <<= b0;
8       for (ulong k=1; k<n; ++k)
9       {
10          ulong xm = (f[k-1] & m ) >> b0;
11          ulong xp = (f[k] & m ) >> b0;
12          if ( xm>xp )  return false;
13      }
14      return true;
15  }
```

The function M is the combination of a mask-out and a shift operation. A routine that sorts according to `b0` and `m` is:

```
1   void
2   counting_sort_core(const ulong * restrict f, ulong n, ulong * restrict g, ulong b0, ulong m)
```

```
// Write to g[] the array f[] sorted wrt. bits b0,...,b0+z-1
//  where z is the number of bits set in m.
// m must contain a single run of bits starting at bit zero.
{
    ulong nb = m + 1;
    m <<= b0;
    ALLOCA(ulong, cv, nb);
    for (ulong k=0; k<nb; ++k)  cv[k] = 0;

    // --- count:
    for (ulong k=0; k<n; ++k)
    {
        ulong x = (f[k] & m ) >> b0;
        ++cv[ x ];
    }

    // --- cumulative sums:
    for (ulong k=1; k<nb; ++k)  cv[k] += cv[k-1];

    // --- reorder:
    ulong k = n;
    while ( k-- )  // backwards ==> stable sort
    {
        ulong fk = f[k];
        ulong x = (fk & m) >> b0;
        --cv[x];
        ulong i = cv[x];
        g[i] = fk;
    }
}
```

```
      Input          Stage 1          Stage 2          Stage 3
                  m = ....11       m = ..11..       m = 11....
                        vv               vv               vv
      111.11          ..1...           11....           ..1...
      ..1...          1111..           1...1.           ..1..1
      .1.1.1          11....           1...11           .1.1.1
      1...1.          .1.1.1           .1.1.1           .1.11.
      1.1111          ..1..1           .1.11.           1...1.
      1111..          1...1.           ..1...           1...11
      ..1..1          .1.11.           ..1..1           1.1111
      .1.11.          111.11           111.11           11....
      1...11          1.1111           1111..           111.11
      11....          1...11           1.1111           1111..
```

Figure 3.11-A: Radix sort of 10 six-bit values when using two-bit masks.

Now we can apply counting sort to a set of bit masks that cover the whole range. Figure 3.11-A shows an example with 10 six-bit values and 3 two-bit masks, starting from the least significant bits. This is the output of the program [FXT: sort/radixsort-demo.cc].

The routine [FXT: `radix_sort()` in sort/radixsort.cc] uses 8-bit masks to sort unsigned (long) integers:

```
void
radix_sort(ulong *f, ulong n)
{
    ulong nb = 8;  // Number of bits sorted with each step
    ulong tnb = BITS_PER_LONG; // Total number of bits

    ulong *fi = f;
    ulong *g = new ulong[n];

    ulong m = (1UL<<nb) - 1;
    for (ulong k=1, b0=0;  b0<tnb;  ++k, b0+=nb)
    {
        counting_sort_core(f, n, g, b0, m);
        swap2(f, g);
    }

    if ( f!=fi )  // result is actually in g[]
    {
        swap2(f, g);
        for (ulong k=0; k<n; ++k)  f[k] = g[k];
    }
```

```
22
23      delete [] g;
24  }
```

There is room for optimization. Using two arrays for counting and combining copying with counting for the next pass where possible will reduce the number of passes almost by a factor of two.

A version of radix sort that starts from the most significant bits is given in [236].

3.12 Searching in unsorted arrays

We give an overview of the search functions for unordered arrays.

3.12.1 Minimal and maximal elements in unsorted arrays

Sorting an array only to determine the minimal (and/or maximal element) is obviously a bad idea because these can be found in linear time in the unsorted data.

Corresponding to all flavors of the sorting routines a `min()` and `max()` routine is supplied. We exemplify by giving the versions that determine the minimum.

Finding the minimum in an unsorted array corresponding to the 'plain' sorting routine: [FXT: sort/minmax.h]

```
1   template <typename Type>
2   Type inline min(const Type *f, ulong n)
3   // Return minimum of array.
4   {
5       Type v = f[0];
6       while ( n-- )  if ( f[n]<v )  v = f[n];
7       return v;
8   }
```

The index version: [FXT: sort/minmaxidx.h]

```
1   template <typename Type>
2   Type idx_min(const Type *f, ulong n, const ulong *x)
3   // Return minimum (value) of array elements
4   //   {f[x[0]], f[x[1]], .. , f[x[n-1]]}
5   {
6       Type v = f[x[0]];
7       while ( n-- )  { if ( f[x[n]]<v )  v = f[x[n]]; }
8       return  v;
9   }
```

The pointer version: [FXT: sort/minmaxptr.h]

```
1   template <typename Type>
2   Type ptr_min(/*const Type *f,*/ ulong n, Type const*const*x)
3   // Return minimum (value) of array elements
4   //   {*x[0], *x[1], .. , *x[n-1]}
5   {
6       Type v = *x[0];
7       while ( n-- )  { if ( *x[n]<v )  v = *x[n]; }
8       return  v;
9   }
```

The comparison function version: [FXT: sort/minmaxfunc.h]

```
1   template <typename Type>
2   Type min(const Type *f, ulong n, int (*cmp)(const Type &, const Type &))
3   // Return minimum (value) of array elements
4   // wrt. to comparison function
5   {
6       Type v = f[0];
7       while ( n-- )  { if ( cmp(f[n],v) < 0 )  v = f[n]; }
8       return  v;
```

```
9   }
```

The comparison function with index version: [FXT: sort/minmaxidxfunc.h]

```
1   template <typename Type>
2   Type idx_min(const Type *f, ulong n, const ulong *x,
3                int (*cmp)(const Type &, const Type &))
4   // Return minimum (value) of array elements
5   //   {f[x[0]], f[x[1]], .. , f[x[n-1]]}
6   // with respect to comparison function cmp()
7   {
8       Type v = f[x[0]];
9       while ( n-- )  { if ( cmp(f[x[n]], v) < 0 )  v = f[x[n]]; }
10      return  v;
11  }
```

The comparison function with pointer version: [FXT: sort/minmaxptrfunc.h]

```
1   template <typename Type>
2   Type ptr_min(/*const Type *f,*/ ulong n, Type const*const*x,
3                int (*cmp)(const Type &, const Type &))
4   // Return minimum (value) of array elements
5   //   {*x[0], *x[1], .. , *x[n-1]}
6   // with respect to comparison function cmp().
7   {
8       Type v = *x[0];
9       while ( n-- )  { if ( cmp(*x[n],v)<0 )  v = *x[n]; }
10      return  v;
11  }
```

3.12.2 Searching for values

To find the first occurrence of a certain value in an unsorted array one can use the routine [FXT: sort/usearch.h]

```
1   template <typename Type>
2   inline ulong first_eq_idx(const Type *f, ulong n, Type v)
3   // Return index of first element == v
4   // Return n if all !=v
5   {
6       ulong k = 0;
7       while ( (k<n) && (f[k]!=v) )  k++;
8       return k;
9   }
```

The functions `first_ne_idx()`, `first_ge_idx()` and `first_le_idx()` find the first occurrence of an element unequal (to `v`), greater or equal and less or equal, respectively.

If the last bit of speed matters, one could (see [166, p.267]) replace the shown code by

```
1   template <typename Type>
2   inline ulong first_eq_idx(const Type *f, ulong n, Type v)
3   // Return index of first element == v
4   // Return n if all !=v
5   {
6       Type s = f[n-1];
7       (Type *)f[n-1] = v;  // guarantee that the search stops
8       ulong k = 0;
9       while ( f[k]!=v )  k++;
10      (Type *)f[n-1] = s;  // restore value
11      if ( (k==n-1) && (v!=s) )  ++k;
12      return k;
13  }
```

The speedup is due to the fact that there is only one branch in the inner loop. The technique is not applicable if the writes to the array '`f[]`' can have any side effects.

Replace the word `first` by last in the names to name the function that detect the last element satisfying the corresponding condition. For example,

```
1   template <typename Type>
2   inline ulong last_ge_idx(const Type *f, ulong n, Type v)
3   // Return index of last element >= v
4   // Return n if all <v
5   {
6       ulong k = n-1;
7       while ( f[k]<v )
8       {
9           k--;
10          if ( 0==k )  return n;
11      }
12      return k;
13  }
```

3.12.3 Counting values

The functions in [FXT: sort/ucount.h] count certain elements in unordered arrays.

```
1   template <typename Type>
2   inline ulong eq_count(const Type *f, ulong n, Type v)
3   // Return number of elements that are ==v
4   {
5       ulong ct = 0;
6       while ( n-- )  if ( f[n]==v )  ++ct;
7       return ct;
8   }
```

As above, replace `eq_` in the function name with `ne_`, `ge_` or `le_` to count the number of occurrences of elements that are unequal (to `v`), greater or equal and less or equal, respectively.

3.12.4 Searching matches

The routines in [FXT: sort/usearchfunc.h] generalize the functions from [FXT: sort/usearch.h]: A supplied function implements the condition imposed. For example,

```
1   template <typename Type>
2   inline ulong first_idx(const Type *f, ulong n, bool (* func)(Type))
3   // Return index of first element for which func() returns true.
4   // Return n if there is no such element.
5   {
6       ulong k = 0;
7       while ( (k<n) && (!func(f[k])) )  k++;
8       return k;
9   }
```

and `last_idx()` that scans beginning from the greatest index.

The functions

```
1   template <typename Type>
2   inline ulong next_idx(const Type *f, ulong n, bool (* func)(Type), ulong k0)
3   // Like first_idx() but start from k0.
4   {
5       ulong k = k0;
6       while ( (k<n) && (!func(f[k])) )  k++;
7       return k;
8   }
```

and `previous_idx()` determine the next matching element in forward or backward direction.

3.12.5 Selecting matches: grep

The routines from [FXT: sort/grep.h] count or select elements from an unordered array for which a condition implemented by a supplied function is true.

The following function counts the matching elements:

```
template <typename Type>
inline ulong count(const Type *f, ulong n, bool (* func)(Type))
// Return number of elements for which func() returns true.
{
    ulong ct = 0;
    for (ulong k=0; k<n; ++k)  if ( func(f[k]) )  ct++;
    return  ct;
}
```

Discard all non-matches using

```
template <typename Type>
inline ulong grep(Type *f, ulong n, bool (* func)(Type))
// Delete elements for which func() returns false.
// Return number of elements kept.
{
    ulong k, j;
    for (k=0,j=0; j<n; ++k,++j)
    {
        f[k] = f[j];
        if ( func(f[j]) )  --k;
    }
    return k;
}
```

Record the values of the matches using

```
template <typename Type>
inline ulong grep(const Type *f, ulong n, bool (* func)(Type), Type *g)
// Make g[] the sequence of values for which func() returns true.
// Return number of 'matching' elements found.
{
    ulong ct = 0;
    for (ulong k=0; k<n; ++k)  if ( func(f[k]) )  g[ct++] = f[k];
    return  ct;
}
```

Finally, recording the indices of the matches can be done with

```
template <typename Type>
inline ulong grep_idx(const Type *f, ulong n, bool (* func)(Type), ulong *x)
// Make x[] the sequence of indices for which func() returns true.
// Return number of 'matching' elements found.
{
    ulong ct = 0;
    for (ulong k=0; k<n; ++k)  if ( func(f[k]) )  x[ct++] = k;
    return  ct;
}
```

Chapter 4

Data structures

We give implementations of simple data structures like stack, ring buffer, queue, double-ended queue (deque), but-array, heap and priority queue. We also describe a finite state engine and the emulation of coroutines.

4.1 Stack (LIFO)

A *stack* (or LIFO for last-in, first-out) is a data structure that supports the operations: *push* to saves an entry and *pop* to retrieve and remove the entry that was entered last. The occasionally useful operation *peek* retrieves the same element as pop but does not remove it. Similarly, *poke* modifies the last entry. An implementation with the option to let the stack grow when necessary is [FXT: `class stack` in ds/stack.h]:

```
template <typename Type>
class stack
{
public:
    Type  *x_;  // data
    ulong  s_;  // size
    ulong  p_;  // stack pointer (position of next write), top entry @ p-1
    ulong  gq_; // grow gq elements if necessary, 0 for "never grow"

public:
    stack(ulong n, ulong growq=0)
    {
        s_ = n;
        x_ = new Type[s_];
        p_ = 0;  // stack is empty
        gq_ = growq;
    }

    ~stack()  { delete [] x_; }

private:
    stack & operator = (const stack &);  // forbidden

public:
    ulong num()  const
    // return number of entries
    {
        return p_;
    }

    ulong push(Type z)
    // return size of stack, zero on stack overflow
    // if gq_ is nonzero the stack grows
    {
        if ( p_ >= s_ )
        {
```

```
37              if ( 0==gq_ )  return 0;  // overflow
38              grow();
39          }
40
41          x_[p_] = z;
42          ++p_;
43
44          return  s_;
45      }
46
47      ulong pop(Type &z)
48      // read top entry and remove it
49      // return number of entries at function start
50      // if empty return zero and z is undefined
51      {
52          ulong ret = p_;
53          if ( 0!=p_ )  { --p_;  z = x_[p_]; }
54          return  ret;
55      }
56
57      ulong poke(Type z)
58      // modify top entry
59      // return number of entries
60      // if empty return zero and nothing is done
61      {
62          if ( 0!=p_ )  x_[p_-1] = z;
63          return p_;
64      }
65
66      ulong peek(Type &z)
67      // read top entry (without removing it)
68      // return number of entries
69      // if empty return zero and z is undefined
70      {
71          if ( 0!=p_ )  z = x_[p_-1];
72          return p_;
73      }
74     [--snip--]
```

The growth routine is implemented as

```
1      [--snip--]
2   private:
3       void grow()
4       {
5           ulong ns = s_ + gq_;  // new size
6           x_ = ReAlloc<Type>(x_, ns, s_);
7           s_ = ns;
8       }
9   };
```

here we used the function that imports the C function `realloc()`.

```
% man realloc
    #include <stdlib.h>

    void *realloc(void *ptr, size_t size);

    realloc() changes the size of the memory block pointed to  by  ptr  to  size
    bytes.   The  contents  will  be unchanged to the minimum of the old and new
    sizes; newly allocated memory will be uninitialized.  If ptr  is  NULL,  the
    call  is  equivalent  to malloc(size); if size is equal to zero, the call is
    equivalent to free(ptr).  Unless ptr is NULL, it must have been returned  by
    an earlier call to malloc(), calloc() or realloc().
```

Usage of the C function `realloc()` can be disabled globally in [FXT: realloc.h]:

```
1   #define USE_C_REALLOC  // comment out to disable use of realloc()
2
3   #ifdef USE_C_REALLOC
4   #include <cstdlib>  // realloc()
5   #endif
6
7   #ifdef USE_C_REALLOC
8   template <typename Type>
9   inline Type *ReAlloc(Type *p, ulong n, ulong /*nold*/)
10  {
11      return (Type *)realloc((void *)p, n*sizeof(Type));
```

```
12  }
13
14  #else
15
16  template <typename Type>
17  inline Type *ReAlloc(Type *p, ulong n, ulong nold)
18  {
19      Type *np = new Type[n];
20      ulong nc = (nold < n ? nold : n);
21      for (ulong k=0; k<nc; ++k)  np[k] = p[k];
22      delete [] p;
23      return np;
24  }
25  #endif
```

```
 push( 1)  1  -  -  -                #=1
 push( 2)  1  2  -  -                #=2
 push( 3)  1  2  3  -                #=3
 push( 4)  1  2  3  4                #=4
 push( 5)  1  2  3  4  5  -  -  -    #=5
 push( 6)  1  2  3  4  5  6  -  -    #=6
 push( 7)  1  2  3  4  5  6  7  -    #=7
 pop== 7   1  2  3  4  5  6  -  -    #=6
 pop== 6   1  2  3  4  5  -  -  -    #=5
 push( 8)  1  2  3  4  5  8  -  -    #=6
 pop== 8   1  2  3  4  5  -  -  -    #=5
 pop== 5   1  2  3  4  -  -  -  -    #=4
 push( 9)  1  2  3  4  9  -  -  -    #=5
 pop== 9   1  2  3  4  -  -  -  -    #=4
 pop== 4   1  2  3  -  -  -  -  -    #=3
 push(10)  1  2  3 10  -  -  -  -    #=4
 pop==10   1  2  3  -  -  -  -  -    #=3
 pop== 3   1  2  -  -  -  -  -  -    #=2
 push(11)  1  2 11  -  -  -  -  -    #=3
 pop==11   1  2  -  -  -  -  -  -    #=2
 pop== 2   1  -  -  -  -  -  -  -    #=1
 push(12)  1 12  -  -  -  -  -  -    #=2
 pop==12   1  -  -  -  -  -  -  -    #=1
 pop== 1   -  -  -  -  -  -  -  -    #=0
 push(13) 13  -  -  -  -  -  -  -    #=1
 pop==13   -  -  -  -  -  -  -  -    #=0
 pop== 0   -  -  -  -  -  -  -  -    #=0
   (stack was empty)
 push(14) 14  -  -  -  -  -  -  -    #=1
 pop==14   -  -  -  -  -  -  -  -    #=0
 pop== 0   -  -  -  -  -  -  -  -    #=0
   (stack was empty)
 push(15) 15  -  -  -  -  -  -  -    #=1
```

Figure 4.1-A: Inserting and retrieving elements with a stack.

A program that shows the working of the stack is [FXT: ds/stack-demo.cc]. An example output where the initial size is 4 and the growths-feature enabled (in steps of 4 elements) is shown in figure 4.1-A.

4.2 Ring buffer

A *ring buffer* is a array plus read- and write operations that wrap around. That is, if the last position of the array is reached writing continues at the begin of the array, thereby erasing the oldest entries. The read operation should start at the oldest entry in the array.

The implementation is [FXT: `class ringbuffer` in ds/ringbuffer.h]:

```
1  template <typename Type>
2  class ringbuffer
3  {
4  public:
5      Type *x_;   // data (ring buffer)
6      ulong s_;   // allocated size (# of elements)
```

```
7       ulong n_;   // current number of entries in buffer
8       ulong wpos_;  // next position to write in buffer
9       ulong fpos_;  // first position to read in buffer
10
11    public:
12      explicit ringbuffer(ulong n)
13      {
14          s_ = n;
15          x_ = new Type[s_];
16          n_ = 0;
17          wpos_ = 0;
18          fpos_ = 0;
19      }
20
21      ~ringbuffer()  { delete [] x_; }
22
23      ulong num()  const  { return n_; }
24
25      void insert(const Type &z)
26      {
27          x_[wpos_] = z;
28          if ( ++wpos_>=s_ )  wpos_ = 0;
29          if ( n_ < s_ )  ++n_;
30          else  fpos_ =  wpos_;
31      }
32
33      ulong read(ulong n, Type &z)  const
34      // read entry n (that is, [(fpos_ + n)%s_])
35      // return 0 if entry #n is last in buffer
36      // else return n+1
37      {
38          if ( n>=n_ )  return 0;
39          ulong j = fpos_ + n;
40          if ( j>=s_ )  j -= s_;
41          z = x_[j];
42          return  n + 1;
43      }
44    };
```

Reading from the ring buffer goes like:

```
ulong k = 0;
Type z;  // type of entry
while ( (k = f.read(k, z)) )
{
    // do something with z
}
```

A program demonstrating the ring buffer is [FXT: ds/ringbuffer-demo.cc], its output is

```
insert( 1)  1            #=1   r=1   w=0
insert( 2)  1  2         #=2   r=2   w=0
insert( 3)  1  2  3      #=3   r=3   w=0
insert( 4)  1  2  3  4   #=4   r=0   w=0
insert( 5)  2  3  4  5   #=4   r=1   w=1
insert( 6)  3  4  5  6   #=4   r=2   w=2
insert( 7)  4  5  6  7   #=4   r=3   w=3
insert( 8)  5  6  7  8   #=4   r=0   w=0
insert( 9)  6  7  8  9   #=4   r=1   w=1
```

Ring buffers can be useful for storing a constant amount of history-data such as for logging purposes. For that purpose one would enhance the `ringbuffer` class so that it uses an additional array of (fixed width) strings. The message to log would be copied into the array and the pointer set accordingly. A read should then just return the pointer to the string.

4.3 Queue (FIFO)

A *queue* (or FIFO for first-in, first-out) is a data structure that supports two operations: *push* saves an entry and *pop* retrieves (and removes) the entry that was entered the longest time ago. The occasionally useful operation *peek* retrieves the same element as pop but does not remove it.

A utility class with the optional feature of growing if necessary is [FXT: `class queue` in ds/queue.h]:

```
template <typename Type>
class queue
{
public:
    Type *x_;   // pointer to data
    ulong s_;   // allocated size (# of elements)
    ulong n_;   // current number of entries in buffer
    ulong wpos_;  // next position to write in buffer
    ulong rpos_;  // next position to read in buffer
    ulong gq_;  // grow gq elements if necessary, 0 for "never grow"

public:
    explicit queue(ulong n, ulong growq=0)
    {
        s_ = n;
        x_ = new Type[s_];
        n_ = 0;
        wpos_ = 0;
        rpos_ = 0;
        gq_ = growq;
    }

    ~queue()  { delete [] x_; }

    ulong num()  const  { return n_; }

    ulong push(const Type &z)
    // Return number of entries.
    // Zero is returned on failure
    //   (i.e. space exhausted and 0==gq_)
    {
        if ( n_ >= s_ )
        {
            if ( 0==gq_ )  return 0;  // growing disabled
            grow();
        }

        x_[wpos_] = z;
        ++wpos_;
        if ( wpos_>=s_ )  wpos_ = 0;

        ++n_;
        return n_;
    }

    ulong peek(Type &z)
    // Return number of entries.
    // if zero is returned the value of z is undefined.
    {
        z = x_[rpos_];
        return n_;
    }

    ulong pop(Type &z)
    // Return number of entries before pop
    // i.e. zero is returned if queue was empty.
    // If zero is returned the value of z is undefined.
    {
        ulong ret = n_;
        if ( 0!=n_ )
        {
            z = x_[rpos_];
            ++rpos_;
            if ( rpos_ >= s_ )  rpos_ = 0;
            --n_;
        }

        return ret;
    }

private:
    void grow()
    {
        ulong ns = s_ + gq_;  // new size
        // move read-position to zero:
```

```
76          rotate_left(x_, s_, rpos_);
77          x_ = ReAlloc<Type>(x_, ns, s_);
78          wpos_ = s_;
79          rpos_ = 0;
80          s_ = ns;
81      }
82  };
```

```
push( 1)  1  -  -  -                #=1   r=0   w=1
push( 2)  1  2  -  -                #=2   r=0   w=2
push( 3)  1  2  3  -                #=3   r=0   w=3
push( 4)  1  2  3  4                #=4   r=0   w=0
push( 5)  1  2  3  4  5  -  -  -    #=5   r=0   w=5
push( 6)  1  2  3  4  5  6  -  -    #=6   r=0   w=6
push( 7)  1  2  3  4  5  6  7  -    #=7   r=0   w=7
pop== 1   -  2  3  4  5  6  7  -    #=6   r=1   w=7
pop== 2   -  -  3  4  5  6  7  -    #=5   r=2   w=7
push( 8)  -  -  3  4  5  6  7  8    #=6   r=2   w=0
pop== 3   -  -  -  4  5  6  7  8    #=5   r=3   w=0
pop== 4   -  -  -  -  5  6  7  8    #=4   r=4   w=0
push( 9)  9  -  -  -  5  6  7  8    #=5   r=4   w=1
pop== 5   9  -  -  -  -  6  7  8    #=4   r=5   w=1
pop== 6   9  -  -  -  -  -  7  8    #=3   r=6   w=1
push(10)  9 10  -  -  -  -  7  8    #=4   r=6   w=2
pop== 7   9 10  -  -  -  -  -  8    #=3   r=7   w=2
pop== 8   9 10  -  -  -  -  -  -    #=2   r=0   w=2
push(11)  9 10 11  -  -  -  -  -    #=3   r=0   w=3
pop== 9   - 10 11  -  -  -  -  -    #=2   r=1   w=3
pop==10   -  - 11  -  -  -  -  -    #=1   r=2   w=3
push(12)  -  - 11 12  -  -  -  -    #=2   r=2   w=4
pop==11   -  -  - 12  -  -  -  -    #=1   r=3   w=4
pop==12   -  -  -  -  -  -  -  -    #=0   r=4   w=4
push(13)  -  -  -  - 13  -  -  -    #=1   r=4   w=5
pop==13   -  -  -  -  -  -  -  -    #=0   r=5   w=5
pop== 0   -  -  -  -  -  -  -  -    #=0   r=5   w=5
  (queue was empty)
push(14)  -  -  -  -  - 14  -  -    #=1   r=5   w=6
pop==14   -  -  -  -  -  -  -  -    #=0   r=6   w=6
pop== 0   -  -  -  -  -  -  -  -    #=0   r=6   w=6
  (queue was empty)
push(15)  -  -  -  -  -  - 15  -    #=1   r=6   w=7
```

Figure 4.3-A: Inserting and retrieving elements with a queue.

Its working is demonstrated by the program [FXT: ds/queue-demo.cc]. An example output where the initial size is 4 and the growths-feature enabled (in steps of 4 elements) is shown in figure 4.3-A. Compare to the corresponding figure for the stack, figure 4.1-A on page 151.

4.4 Deque (double-ended queue)

A *deque* (for *double-ended queue*) combines the data structures stack and queue: insertion and deletion is possible both at the first- and the last position, all in time-$O(1)$. An implementation with the option to let the stack grow when necessary is [FXT: `class deque` in ds/deque.h]

```
1   template <typename Type>
2   class deque
3   {
4   public:
5       Type *x_;   // data (ring buffer)
6       ulong s_;   // allocated size (# of elements)
7       ulong n_;   // current number of entries in buffer
8       ulong fpos_;  // position of first element in buffer
9       // insert_first() will write to (fpos-1)%n
10      ulong lpos_;  // position of last element in buffer plus one
11      // insert_last() will write to  lpos,  n==(lpos-fpos) (mod s)
12      // entries are at [fpos, ... , lpos-1]  (range may be empty)
```

```
13
14      ulong  gq_; // grow gq elements if necessary, 0 for "never grow"
15
16  public:
17      explicit deque(ulong n, ulong growq=0)
18      {
19          s_ = n;
20          x_ = new Type[s_];
21          n_ = 0;
22          fpos_ = 0;
23          lpos_ = 0;
24          gq_ = growq;
25      }
26
27      ~deque()  { delete [] x_; }
28
29      ulong num()  const  { return n_; }
30
31      ulong insert_first(const Type &z)
32      // returns number of entries after insertion
33      // zero is returned on failure
34      //   (i.e. space exhausted and 0==gq_)
35      {
36          if ( n_ >= s_ )
37          {
38              if ( 0==gq_ )  return 0;  // growing disabled
39              grow();
40          }
41
42          --fpos_;
43          if ( fpos_ == -1UL )  fpos_ = s_ - 1;
44          x_[fpos_] = z;
45          ++n_;
46          return  n_;
47      }
48
49
50      ulong insert_last(const Type &z)
51      // returns number of entries after insertion
52      // zero is returned on failure
53      //   (i.e. space exhausted and 0==gq_)
54      {
55          if ( n_ >= s_ )
56          {
57              if ( 0==gq_ )  return 0;  // growing disabled
58              grow();
59          }
60
61          x_[lpos_] = z;
62          ++lpos_;
63          if ( lpos_>=s_ )  lpos_ = 0;
64          ++n_;
65          return  n_;
66      }
67
68
69      ulong extract_first(Type & z)
70      // return number of elements before extract
71      // return 0 if extract on empty deque was attempted
72      {
73          if ( 0==n_ )  return 0;
74          z = x_[fpos_];
75          ++fpos_;
76          if ( fpos_ >= s_ )  fpos_ = 0;
77          --n_;
78          return  n_ + 1;
79      }
80
81      ulong extract_last(Type & z)
82      // return number of elements before extract
83      // return 0 if extract on empty deque was attempted
84      {
85          if ( 0==n_ )  return 0;
86          --lpos_;
87          if ( lpos_ == -1UL )  lpos_ = s_ - 1;
88          z = x_[lpos_];
89          --n_;
```

```
90          return  n_ + 1;
91      }
92
93
94      ulong read_first(Type & z)  const
95      // read (but don't remove) first entry
96      // return number of elements (i.e. on error return zero)
97      {
98          if ( 0==n_ )  return 0;
99          z = x_[fpos_];
100         return n_;
101     }
102
103     ulong read_last(Type & z)  const
104     // read (but don't remove) last entry
105     // return number of elements (i.e. on error return zero)
106     {
107         return  read(n_-1, z);  // ok for n_==0
108     }
109
110     ulong read(ulong k, Type & z)  const
111     // read entry k (that is, [(fpos_ + k)%s_])
112     // return 0 if k>=n_
113     // else return k+1
114     {
115         if ( k>=n_ )  return 0;
116         ulong j = fpos_ + k;
117         if ( j>=s_ )  j -= s_;
118         z = x_[j];
119         return  k + 1;
120     }
121
122 private:
123     void grow()
124     {
125         ulong ns = s_ + gq_;  // new size
126         // Move read-position to zero:
127         rotate_left(x_, s_, fpos_);
128         x_ = ReAlloc<Type>(x_, ns, s_);
129         fpos_ = 0;
130         lpos_ = n_;
131         s_ = ns;
132     }
133 };
```

```
insert_first( 1)        1
insert_last(51)         1 51
insert_first( 2)        2  1 51
insert_last(52)         2  1 51 52
insert_first( 3)        3  2  1 51 52
insert_last(53)         3  2  1 51 52 53
extract_first()= 3      2  1 51 52 53
extract_last()= 53      2  1 51 52
insert_first( 4)        4  2  1 51 52
insert_last(54)         4  2  1 51 52 54
extract_first()= 4      2  1 51 52 54
extract_last()= 54      2  1 51 52
extract_first()= 2      1 51 52
extract_last()= 52      1 51
extract_first()= 1     51
extract_last()= 51
insert_first( 5)        5
insert_last(55)         5 55
extract_first()= 5     55
extract_last()= 55
extract_first()= (deque is empty)
extract_last()=  (deque is empty)
insert_first( 7)        7
insert_last(57)         7 57
```

Figure 4.4-A: Inserting and retrieving elements with a queue.

Its working is shown in figure 4.4-A which was created with the program [FXT: ds/deque-demo.cc].

4.5 Heap and priority queue

A *binary heap* is a binary tree where the left and right children are smaller than or equal to their parent node. The following function determines whether a given array is a heap [FXT: ds/heap.h]:

```
template <typename Type>
ulong test_heap(const Type *x, ulong n)
// return 0 if x[] has heap property
// else index of node found to be bigger than its parent
{
    const Type *p = x - 1;
    for (ulong k=n; k>1; --k)
    {
        ulong t = (k>>1); // parent(k)
        if ( p[t]<p[k] )  return k;
    }
    return 0;  // has heap property
}
```

It turns out that a unordered array of size n can be reordered to a heap in $O(n)$ time, see [95, p.145]. The routine

```
template <typename Type>
void build_heap(Type *x, ulong n)
// reorder data to a heap
{
    for (ulong j=(n>>1); j>0; --j)  heapify1(x-1, n, j);
}
```

does the trick. It uses the following subroutine that runs in time $O(\log n)$:

```
template <typename Type>
void heapify1(Type *z, ulong n, ulong k)
// subject to the condition that the trees below the children of node
// k are heaps, move the element z[k] (down) until the tree below node
// k is a heap.
//
// data expected in z[1..n] (!)
{
    ulong m = k;
    ulong l = (k<<1); // left(k);
    if ( (l <= n) && (z[l] > z[k]) )  m = l;
    ulong r = (k<<1) + 1; // right(k);
    if ( (r <= n) && (z[r] > z[m]) )  m = r;
    if ( m != k )
    {
        swap2(z[k], z[m]);
        heapify1(z, n, m);
    }
}
```

Heaps are useful to build a so-called *priority queue*. This is a data structure that supports insertion of an element and extraction of the maximal element it contains both in an efficient manner. A priority queue can be used to 'schedule' a certain event for a given point in time and return the next pending event.

A new element can be inserted into a heap in $O(\log(n))$ time by appending it and moving it towards the root (that is, the first element) as necessary:

```
bool heap_insert(Type *x, ulong n, ulong s, Type t)
// with x[] a heap of current size n
// and max size s (i.e. space for s elements allocated)
// insert t and restore heap-property.
//
// Return true if successful
// else (i.e. space exhausted) false
//
// Takes time \log(n)
{
    if ( n > s )  return false;
    ++n;
    Type *x1 = x - 1;
    ulong j = n;
```

```
15      while ( j > 1 )
16      {
17          ulong k = (j>>1);  // k==parent(j)
18          if ( x1[k] >= t )  break;
19          x1[j] = x1[k];
20          j = k;
21      }
22      x1[j] = t;
23      return true;
24  }
```

Similarly, the maximal element can be removed in time $O(\log(n))$:

```
1   template <typename Type>
2   Type heap_extract_max(Type *x, ulong n)
3   // Return maximal element of heap and
4   // restore heap structure.
5   // Return value is undefined for 0==n
6   {
7       Type m = x[0];
8       if ( 0 != n )
9       {
10          Type *x1 = x - 1;
11          x1[1] = x1[n];
12          --n;
13          heapify1(x1, n, 1);
14      }
15      // else error
16      return m;
17  }
```

Two modifications seem appropriate: firstly, replacement of `extract_max()` by a `extract_next()`, leaving it as an (compile time) option whether to extract the minimal or the maximal element. This is achieved by changing the comparison operators at a few strategic places so that the heap is built either as described above or with its minimum as first element:

```
1   #if  1
2   // next() is the one with the smallest key
3   // i.e.  extract_next()  is  extract_min()
4   #define _CMP_ <
5   #define _CMPEQ_ <=
6   #else
7   // next() is the one with the biggest key
8   // i.e.  extract_next()  is  extract_max()
9   #define _CMP_ >
10  #define _CMPEQ_ >=
11  #endif
```

Secondly, augmenting the elements used by a event description that can be freely defined. [FXT: `class priority_queue` in ds/priorityqueue.h]:

```
1   template <typename Type1, typename Type2>
2   class priority_queue
3   {
4   public:
5       Type1 *t1_;  // time:   t1[1..s]  one-based array!
6       Type2 *e1_;  // events: e1[1..s]  one-based array!
7       ulong s_;   // allocated size (# of elements)
8       ulong n_;   // current number of events
9       ulong gq_;  // grow gq elements if necessary, 0 for "never grow"
10
11
12  public:
13      priority_queue(ulong n, ulong growq=0)
14      {
15          s_ = n;
16          t1_ = new Type1[s_] - 1;
17          e1_ = new Type2[s_] - 1;
18
19          n_ = 0;
20          gq_ = growq;
21      }
22
23      ~priority_queue()
24      {
25          delete [] (t1_+1);
26          delete [] (e1_+1);
```

```
27      }
28
29      ulong num()  const  { return n_; }
30
31      Type1 get_next_t()  const  { return t1_[1]; }
32      Type2 get_next_e()  const  { return e1_[1]; }
33
34      Type1 get_next(Type2 &e)  const
35      // No check if empty
36      {
37          e = e1_[1];
38          return t1_[1];
39      }
40
41      Type1 extract_next(Type2 &e)
42      {
43          Type1 m = get_next(e);
44          if ( 0 != n_ )
45          {
46              t1_[1] = t1_[n_];  e1_[1] = e1_[n_];
47              --n_;
48              heapify1(1);
49          }
50          // else error
51          return m;
52      }
53
54      bool insert(const Type1 &t, const Type2 &e)
55      // Insert event e at time t
56      // Return true if successful
57      // else (i.e. space exhausted and 0==gq_) false
58      {
59          if ( n_ >= s_ )
60          {
61              if ( 0==gq_ )  return false;  // growing disabled
62              grow();
63          }
64
65          ++n_;
66          ulong j = n_;
67          while ( j > 1 )
68          {
69              ulong k = (j>>1);  // k==parent(j)
70              if ( t1_[k] _CMPEQ_ t )  break;
71              t1_[j] = t1_[k];  e1_[j] = e1_[k];
72              j = k;
73          }
74          t1_[j] = t;
75          e1_[j] = e;
76
77          return true;
78      }
```

The member function `reschedule_next()` is more efficient than the sequence `extract_next(); insert();`, as it calls `heapify1()` only once:

```
1       void reschedule_next(Type1 t)
2       {
3           t1_[1] = t;
4           heapify1(1);
5       }
```

The `heapify1()` function is tail recursive, so we make it iterative:

```
1   private:
2       void heapify1(ulong k)
3       {
4           ulong m = k;
5
6       hstart:
7           ulong l = (k<<1); // left(k);
8           ulong r = l + 1; // right(k);
9           if ( (l <= n_) && (t1_[l] _CMP_ t1_[k]) )  m = l;
10          if ( (r <= n_) && (t1_[r] _CMP_ t1_[m]) )  m = r;
11
12          if ( m != k )
13          {
14              swap2(t1_[k], t1_[m]);  swap2(e1_[k], e1_[m]);
```

```
15  //            heapify1(m);
16              k = m;
17              goto hstart;  // tail recursion
18          }
19      }
20      [--snip--]
```

The second argument of the constructor determines the number of elements added in case of growth, it is disabled (equals zero) by default.

```
1       [--snip--]
2   private:
3       void grow()
4       {
5           ulong ns = s_ + gq_;  // new size
6           t1_ = ReAlloc<Type1>(t1_+1, ns, s_) - 1;
7           e1_ = ReAlloc<Type2>(e1_+1, ns, s_) - 1;
8           s_ = ns;
9       }
10  };
```

The `ReAlloc()` routine is described in section 4.1 on page 149.

```
Inserting into piority_queue:          Extracting from piority_queue:
 # :  event  @          time            # :  event  @          time
 0:       A  @      0.840188            9:       F  @      0.197551
 1:       B  @      0.394383            8:       I  @      0.277775
 2:       C  @      0.783099            7:       G  @      0.335223
 3:       D  @       0.79844            6:       B  @      0.394383
 4:       E  @      0.911647            5:       J  @       0.55397
 5:       F  @      0.197551            4:       H  @       0.76823
 6:       G  @      0.335223            3:       C  @      0.783099
 7:       H  @       0.76823            2:       D  @       0.79844
 8:       I  @      0.277775            1:       A  @      0.840188
 9:       J  @       0.55397            0:       E  @      0.911647
```

Figure 4.5-A: Insertion of events labeled '`A`', '`B`', ..., '`J`' scheduled for random times into a priority queue (left) and subsequent extraction (right).

The program [FXT: ds/priorityqueue-demo.cc] inserts events at random times $0 \le t < 1$, then extracts all of them. It gives the output shown in figure 4.5-A. A more typical usage would intermix the insertions and extractions.

4.6 Bit-array

The use of *bit-arrays* should be obvious: an array of tag values (like 'seen' versus 'unseen') where all standard data types would be a waste of space. Besides reading and writing individual bits one should implement a convenient search for the next set (or cleared) bit.

The class [FXT: `class bitarray` in ds/bitarray.h] is used, for example, for lists of small primes [FXT: mod/primes.cc], for in-place transposition routines [FXT: aux2/transpose.h] (see section 2.3 on page 98) and several operations on permutations (see section 2.11.8 on page 119).

```
1   class bitarray
2   // Bit-array class mostly for use as memory saving array of boolean values.
3   // Valid index is 0...nb_-1 (as usual in C arrays).
4   {
5   public:
6       ulong *f_;   // bit bucket
7       ulong n_;    // number of bits
8       ulong nfw_;  // number of words where all bits are used, may be zero
9       ulong mp_;   // mask for partially used word if there is one, else zero
10      // (ones are at the positions of the _unused_ bits)
11      bool myfq_;  // whether f[] was allocated by class
12    [--snip--]
```

The constructor allocates memory by default. If the second argument is nonzero it must point to an accessible memory range:

```
1       bitarray(ulong nbits, ulong *f=0)
2       // nbits must be nonzero
3       {
4           ulong nw = ctor_core(nbits);
5           if ( f!=0 )
6           {
7               f_ = (ulong *)f;
8               myfq_ = false;
9           }
10          else
11          {
12              f_ = new ulong[nw];
13              myfq_ = true;
14          }
15      }
```

The public methods are

```
1   // operations on bit n:
2       ulong test(ulong n)  const // Test whether n-th bit set
3       void set(ulong n)          // Set n-th bit
4       void clear(ulong n)        // Clear n-th bit
5       void change(ulong n)       // Toggle n-th bit
6       ulong test_set(ulong n)    // Test whether n-th bit is set and set it
7       ulong test_clear(ulong n)  // Test whether n-th bit is set and clear it
8       ulong test_change(ulong n) // Test whether n-th bit is set and toggle it
9
10  // Operations on all bits:
11      void clear_all()           // Clear all bits
12      void set_all()             // Set all bits
13      int all_set_q()  const;    // Return whether all bits are set
14      int all_clear_q()  const;  // Return whether all bits are clear
15
16  // Scanning the array:
17      // Note: the given index n is included in the search
18      ulong next_set_idx(ulong n)  const   // Return index of next set or value beyond end
19      ulong next_clear_idx(ulong n)  const // Return index of next clear or value beyond end
```

Combined operations like 'test-and-set-bit', 'test-and-clear-bit', 'test-and-change-bit' are often needed in applications that use bit-arrays. This is underlined by the fact that modern CPUs typically have instructions implementing these operations.

The class does not supply overloading of the array-index operator `[]` because the writing variant would cause a performance penalty. One might want to add 'sparse'-versions of the scan functions (`next_set_idx()` and `next_clear_idx()`) for large bit-arrays with only few bits set or unset.

On the AMD64 architecture the corresponding CPU instructions are used [FXT: bits/bitasm-amd64.h]:

```
1   static inline ulong asm_bts(ulong *f, ulong i)
2   // Bit Test and Set
3   {
4       ulong ret;
5       asm ( "btsq %2, %1 \n"
6             "sbbq %0, %0"
7             : "=r" (ret)
8             : "m" (*f), "r" (i) );
9       return ret;
10  }
```

If no specialized CPU instructions are available the following two macros are used:

```
1   #define  DIVMOD(n, d, bm) \
2   ulong d = n / BITS_PER_LONG; \
3   ulong bm = 1UL << (n % BITS_PER_LONG);
```

```
1   #define  DIVMOD_TEST(n, d, bm) \
2   ulong d = n / BITS_PER_LONG; \
3   ulong bm = 1UL << (n % BITS_PER_LONG); \
4   ulong t = bm & f_[d];
```

An example:

```
1       ulong test_set(ulong n)  // Test whether n-th bit is set and set it.
2       {
3   #ifdef  BITS_USE_ASM
4           return asm_bts(f_, n);
5   #else
6           DIVMOD_TEST(n, d, bm);
7           f_[d] |= bm;
8           return  t;
9   #endif
10      }
```

Performance is still good in that case as the modulo operation and division by `BITS_PER_LONG` (a power of two) are replaced with cheap (bit-and and shift) operations. On the machine described in appendix B on page 913 both versions give practically identical performance.

How out of bounds are handled can be defined at the beginning of the header file:

```
#define CHECK  0  // define to disable check of out of bounds access
//#define CHECK  1  // define to handle out of bounds access
//#define CHECK  2  // define to fail with out of bounds access
```

4.7 Finite-state machines

A *finite-state machine* (FSM) (or *state-engine*) in its simplest form can be described as a program that has finite set of valid states and for each state a certain action is taken. In C-syntax:

```
1   void state_engine(int state)
2   {
3       while ( state != end )
4       {
5           // valid states are: st1 ... stn (and end)
6           switch ( state )
7           {
8           case st1:  state = func1();  break;
9           case st2:  state = func2();  break;
10          case st3:  state = func3();  break;
11          [--snip--]
12          case stn:  state = funcn();  break;
13          default:   blue_smoke();  // invalid state
14          }
15      }
16  }
```

```
int main()
{
    // initialize:
    int state = start;
    state_engine( state );
    return 0;
}
```

As an example we show a state engine that transforms a linear coordinate t into the corresponding pair (x and y) of coordinates of Hilbert's space-filling curve.

Apart from two bits of internal state the FSM processes at each step two bits of input. The array `htab[]` serves as lookup table for the next state plus two bits of the result.

The function ([FXT: `lin2hilbert()` in bits/hilbert.cc]) implements a FSM as suggested in [31], item 115:

```
1   void
2   lin2hilbert(ulong t, ulong &x, ulong &y)
3   // Transform linear coordinate t to Hilbert x and y
4   {
5       ulong xv = 0, yv = 0;
6       ulong c01 = (0<<2);  // (2<<2) for transposed output (swapped x, y)
7       for (ulong i=0; i<(BITS_PER_LONG/2); ++i)
8       {
9           ulong abi = t >> (BITS_PER_LONG-2);
10          t <<= 2;
11
12          ulong st = htab[ (c01<<2) | abi ];
13          c01 = st & 3;
```

```
14
15          yv <<= 1;
16          yv |= ((st>>2) & 1);
17          xv <<= 1;
18          xv |= (st>>3);
19      }
20      x = xv;  y = yv;
21  }
```

OLD				NEW				NEW				OLD			
C	C	A	B	X	Y	C	C	C	C	X	Y	A	B	C	C
0	1	I	I	I	I	0	1	0	1	I	I	I	I	0	1
0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	0
0	0	0	1	0	1	0	0	0	0	0	1	0	1	0	0
0	0	1	0	1	1	0	0	0	0	1	0	1	1	0	1
0	0	1	1	1	0	0	1	0	0	1	1	1	0	0	0
0	1	0	0	1	1	1	1	0	1	0	0	1	0	0	1
0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
0	1	1	0	0	0	0	1	0	1	1	0	1	1	0	0
0	1	1	1	1	0	0	0	0	1	1	1	0	0	1	1
1	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
1	0	0	1	1	0	1	0	1	0	0	1	1	1	1	1
1	0	1	0	1	1	1	0	1	0	1	0	0	1	1	0
1	0	1	1	0	1	1	1	1	0	1	1	1	0	1	0
1	1	0	0	1	1	0	1	1	1	0	0	1	0	1	1
1	1	0	1	1	0	1	1	1	1	0	1	1	1	1	0
1	1	1	0	0	0	1	1	1	1	1	0	0	1	1	1
1	1	1	1	0	1	1	0	1	1	1	1	0	0	0	1

Figure 4.7-A: The original table from [31] for the finite state machine for the 2-dimensional Hilbert curve (left). All sixteen 4-bit words appear in both the '`OLD`' and the '`NEW`' column. So the algorithm is invertible. Swap the columns and sort numerically to obtain the two columns at the right, the table for the inverse function.

The table used is defined (see figure 4.7-A) as

```
1   static const ulong htab[] = {
2   #define HT(xi,yi,c0,c1) ((xi<<3)+(yi<<2)+(c0<<1)+(c1))
3       // index == HT(c0,c1,ai,bi)
4       HT( 0, 0,  1, 0 ),
5       HT( 0, 1,  0, 0 ),
6       HT( 1, 1,  0, 0 ),
7       HT( 1, 0,  0, 1 ),
8     [--snip--]
9       HT( 0, 0,  1, 1 ),
10      HT( 0, 1,  1, 0 )
11  };
```

As indicated in the code the table maps every four bits `c0,c1,ai,bi` to four bits `xi,yi,c0,c1`. The table for the inverse function (again, see figure 4.7-A) is

```
1   static const ulong ihtab[] = {
2   #define IHT(ai,bi,c0,c1) ((ai<<3)+(bi<<2)+(c0<<1)+(c1))
3       // index == HT(c0,c1,xi,yi)
4       IHT( 0, 0,  1, 0 ),
5       IHT( 0, 1,  0, 0 ),
6       IHT( 1, 1,  0, 1 ),
7       IHT( 1, 0,  0, 0 ),
8     [--snip--]
9       IHT( 0, 1,  1, 1 ),
10      IHT( 0, 0,  0, 1 )
11  };
```

The words have to be processed backwards:

```
1   ulong
2   hilbert2lin(ulong x, ulong y)
3   // Transform Hilbert x and y to linear coordinate t
4   {
5       ulong t = 0;
6       ulong c01 = 0;
7       for (ulong i=0; i<(BITS_PER_LONG/2); ++i)
8       {
9           t <<= 2;
```

```
10          ulong xi = x >> (BITS_PER_LONG/2-1);
11          xi &= 1;
12          ulong yi = y >> (BITS_PER_LONG/2-1);
13          yi &= 1;
14          ulong xyi = (xi<<1) | yi;
15          x <<= 1;
16          y <<= 1;
17
18          ulong st = ihtab[ (c01<<2) | xyi ];
19          c01 = st & 3;
20
21          t |= (st>>2);
22      }
23
24      return t;
25  }
```

A method to compute the direction (left, right, up or down) at the n-th move of the Hilbert curve is given in section 1.20 on page 56. The computation of a function whose series coefficients are ± 1 and $\pm i$ according to the Hilbert curve is described in section 36.9 on page 734.

4.8 Emulation of coroutines

Sometimes it is possible to find an recursive algorithm for solving some problem that is not easily solved iteratively. However the recursive implementations might produce the results in midst of its calling graph. When a utility class providing the results one by one with some `next` call is required there is an apparent problem (with most programming languages): there is only one stack available for function calls. We do not have 'offline' functions or *coroutines*.

As an example consider the following recursive code (given by Glenn Rhoads)

```
1   int n = 4;
2   int a[n+1], q[2*n+1];
3
4   void paren(long i, long s)
5   {
6       long k, t;
7       if ( i<n )
8       {
9           for (k=0; k<=i-s;  ++k)
10          {
11              a[i-1] = k;
12              t = s + a[i-1];
13              q[t + i] = '(';
14              paren(i + 1, t);  // recursion
15              q[t + i] = ')';
16          }
17      }
18      else
19      {
20          a[i-1] = n - s;
21          Visit();  // next set of parens available
22      }
23  }
24
25  int main()
26  {
27      paren(0, 0);
28      return 0;
29  }
```

It generates following output (the different ways to group four pairs of parenthesis):

```
(((())))
((()()))
((())())
((()))()
(()(()))
(()()())
(()())()
(())(())
```

```
((())()()
()((()))
()(()())
()(())()
()()(())
()()()()
```

A reasonable way to create coroutines is to rewrite the function as a state engine and use a class [FXT: `class coroutine` in ds/coroutine.h] that provides two stacks, one for local variables (including arguments) and one for the state of the function:

```
template <typename Type>
class coroutine
{
public:
    ulong *t_;  // sTate stack
    ulong tp_;  // sTate stack Pointer
    Type  *d_;  // Data stack
    ulong dp_;  // Data stack Pointer
    ulong  n_;  // size of stacks

public:
    coroutine(ulong maxdepth)
    {
        n_ = maxdepth;
        t_ = new ulong[n_];
        d_ = new Type[n_];
        init();
    }

    ~coroutine()
    {
        delete [] d_;
        delete [] t_;
    }

    void init()  { dp_=0; tp_=0; }

    // --- state stack:
    void stpush(ulong x)   { chk_tp();  t_[tp_] = x;  ++tp_; }
    ulong stpeek()  const  { chk_tp(1);  return  t_[tp_-1]; }
    void stpoke(ulong x)   { chk_tp(1);  t_[tp_-1] = x; }
    void stpop(ulong ct=1)  { tp_-=ct; }
    void stnext()  { chk_tp(1);  ++t_[tp_-1]; }
    void stnext(ulong x)  { chk_tp(1);  t_[tp_-1] = x; }
    bool more()  const  { return (tp_!=0); }

    // --- stack for variables and args:
    void push(Type x)  {  chk_dp();  d_[dp_] = x;  ++dp_; }
    Type &peek()  { chk_dp(1); return d_[dp_-1]; }
    void pop(ulong ct=1) { dp_-=ct; }


private:
    coroutine & operator = (const coroutine &);  // forbidden
};
```

The functions `chk_tp()` and`chk_dp()` either check for stack overflow or do nothing, as defined by

```
#define CHECK_STACKS  // define to detect stack overflow
```

Rewriting the function in question (as part of a utility class, given in [FXT: ds/coroutine-paren-demo.cc]) only requires the understanding of the language, not of the algorithm. The process is straightforward but needs a bit of concentration:

```
int
paren::next_recursion()
{
 redo:
    vars &V = cr_->peek();
    int &i=V.i,  &s=V.s; // args
    int &k=V.k,  &t=V.t; // locals

 loop:
    switch ( cr_->stpeek() )
    {
    case 0:
```

```
13          if ( i>=n )
14          {
15              x[i-1] = n - s;
16              cr_->stnext( ST_RETURN );  return 1;
17          }
18          cr_->stnext();
19
20      case 1:
21          // loop end ?
22          if ( k>i-s )  { break; }  // shortcut: nothing to do at end
23
24          cr_->stnext();
25
26      case 2: // start of loop body
27          {
28              long ii = i - 1;
29              x[ii] = k;
30              t = s + k;
31
32              str[t+i] = '('; // OPEN_CHAR;
33
34              cr_->stnext();
35              CORO_CALL( i+1, t, 0, 0 );
36              goto redo;
37          }
38
39      case 3:
40          str[t+i] = ')'; // CLOSE_CHAR;
41          ++k;
42
43          // loop end ?
44          if ( k>i-s )  { break; }  // shortcut: nothing to do at end
45
46          cr_->stpoke(2);  goto loop; // shortcut: back to loop body
47
48      default: ;
49      }
50
51      CORO_RETURN();
52
53      if ( cr_->more() )  goto redo;
54
55      return 0;
56  }
```

The following `#define` was used:

```
1   //           args=(i, s)  (k, t)=locals
2   #define CORO_CALL(vi, vs, vk, vt) \
3    { vars V_; V_.i=vi; V_.s=vs; V_.k=vk; V_.t=vt; \
4      cr_->push(V_);  cr_->stpush(ST_INIT); }
```

The constructor pushes the needed variables and parameters on the data stack and the initial state on the state stack:

```
1   paren::paren(int nn)
2   {
3       n = (nn>0 ? nn : 1);
4       x = new int[n+1];
5       ++x;
6
7       str = new char[2*n+1];
8       int i = 0;
9       for (  ; i<n  ; ++i)  str[i] = '('; // OPEN_CHAR;
10      for (  ; i<2*n; ++i)  str[i] = ')'; // CLOSE_CHAR;
11      str[2*n] = 0;
12
13      cr_ = new coroutine<vars>(n+1);
14
15      //         i, s, k, t
16      CORO_CALL( 0, 0, 0, 0 );
17
18      idx = 0;
19      q = next_recursion();
20  }
```

The method `next()` of the paren class lets the offline function advance until the next result is available:

```
1   int paren::next()
2   {
3       if ( 0==q )  return 0;
```

```
4       else
5       {
6           q = next_recursion();
7           return  ( q ? ++idx : 0 );
8       }
9   }
```

Performance-wise the `coroutine`-rewritten functions are close to the original: state engines are fast and the stack operations are cheap.

The shown method can also be applied when the recursive algorithm consists of more than one function by merging the functions into one state engine.

Further, investigating the contents of the data stack can be of help in the search of a iterative solution.

The code of the converted functions is admittedly ugly but without a generic support of coroutines it seems hard to du much better. One could also use threads for the emulation of coroutines. This approach, however, adds portability problems.

An iterative algorithm for the generation of valid pairings or parenthesis is given in chapter 13 on page 313.

Part II

Combinatorial generation

Chapter 5

Conventions and considerations

We give algorithms for the generation of all combinatorial objects of certain types such as for the generation of combinations, compositions, subsets, permutations, integer partitions, set partitions, restricted growth strings and necklaces. Finally, we give some constructions for Hadamard and conference matrices. Several (more esoteric) combinatorial objects that are found via searching in directed graphs are presented in chapter 19.

The routines are useful in situations where an exhaustive search over all configurations of a certain kind is needed. Combinatorial algorithms are also fundamental with many programming problems and finally they can simply be fun!

5.1 About representations and orders

For a set of n elements we will always take either $\{0, 1, \ldots, n-1\}$ or $\{1, 2, \ldots, n\}$. Our convention for the set notation is to always start with the smallest element. Often there is more than one useful way to represent a combinatorial object. For example the subset $\{1, 4, 6\}$ of the set $\{0, 1, 2, 3, 4, 5, 6\}$ can also be written as a *delta set* `[0100101]`. Some sources use the term *bit string*. We often write dots for zeros for readability: `[.1..1.1]`. Note that in the delta set we put the first element to the left side (*array notation*), this is in contrast to the usual way to print binary numbers, where the least significant bit (bit number zero) is shown on the right side.

For most objects we will give an algorithm for generation in *lexicographic* (or simply *lex*) order. In lexicographic order a string $X = [x_0, x_1, \ldots]$ precedes another string $Y = [y_0, y_1, \ldots]$ if for the smallest index k where the strings differ we have $x_k < y_k$. Further, the string X precedes $X.W$ (the concatenation of X with W) for any nonempty string W. The ordering obtained by reversing the lexicographic order is sometimes called *relex* order. The *co-lexicographic* (or simply *colex*) order is obtained by the lex order of the reversed strings. The order sometimes depends on the representation that is used, see figure 8.1-A on page 201.

In a *minimal-change* order the amount of change between successive objects is the least possible. Such an order is also called a (combinatorial) *Gray code*. There is in general more than one such order. Often one can impose even stricter conditions, like that (with permutations) the changes are between neighboring positions. The corresponding order is a *strong minimal-change order*. A very readable survey of Gray codes in given in [262], see also [230].

5.2 Ranking, unranking, and counting

For a particular ordering of combinatorial objects (say, lexicographic order for permutations) one can ask which position in the list a given object has. An algorithm for finding the position is called a *ranking* algorithm. An method to determine the object, given its position, is called an *unranking* algorithm.

Given both ranking and unranking methods one can compute the successor of a given object by computing its rank r and unranking $r + 1$. While this method is usually slow the idea can be used to find more efficient algorithms for computing the successor. In addition the idea often suggests interesting orderings for combinatorial objects.

We sometimes give ranking or unranking methods for numbers in special forms such as factorial representations for permutations. Ranking and unranking methods are implicit in generation algorithms based on mixed radix counting given in section 10.8 on page 258.

A simple but surprisingly powerful way for the discovery of isomorphisms (one-to-one correspondences) between combinatorial objects is counting them. If the sequences of numbers of two kinds of objects are identical chances are good to find a conversion routine between the corresponding objects. For example, there are 2^n permutations of n elements such that no element lies more than one position to the right of its original position. From this observation an algorithm for generating these permutations via binary counting can be found, see section 10.15.2 on page 286.

The representation of combinatorial objects as restricted growth strings (as shown in section 13.2 on page 315) follows from the same idea. The resulting generation methods can be very fast and flexible.

Recursive relations for the number of objects often lead to recursive generation algorithms (see, for example, section 6.3 on page 181 and section 11.2.1 on page 291).

5.3 Characteristics of the algorithms

In almost all cases we produce the combinatorial objects on by one. Let n be the size of the object. The successor (with respect to the specified order) is computed from the object itself and additional data whose size is less than a constant multiple of n.

Let B be the total number of combinatorial objects under consideration. Sometimes the cost of a successor computation is proportional to n. Then the total cost for generating all objects is proportional to $n \cdot B$.

If the successor computation takes a fixed number of operations (independent of the object size) then we say the algorithm is $O(1)$. If so, there can be no loop in the implementation, we say the algorithm is *loopless*. The total cost for all objects then is $c \cdot B$ for some c independent of the object size. A loopless algorithm can only exist is the amount of change between successive object is bounded by a constant that does not depend on the object size. Natural candidates for loopless algorithms are Gray codes.

In many cases the cost of computing all objects is also $c \cdot B$ while the computation of the successor does involve a loop. As an example consider incrementing in binary using arrays: in half of the cases just the lowest bit changes, for half of the remaining cases just two bits change, and so on. The total cost is $B \cdot (1 + \frac{1}{2}(1 + \frac{1}{2}(\cdots))) = 2 \cdot B$, independent of the number of bits used. So the total cost is as in the loopless case while the successor computation can be expensive in some cases. Algorithms with this characteristic are called *constant amortized time* (or CAT). Often CAT algorithms are faster than loopless algorithms, typically when their structure is more simple.

5.4 Optimization techniques

Let `x` be an array of n elements. The loop

```
ulong k = 0;
while ( (k<n) && (x[k]!=0) )  ++k;  // find first zero
```

can be replaced by

```
ulong k = 0;
while ( x[k]!=0 )  ++k;  // find first zero
```

if a single *sentinel* element `x[n]=0` is appended to the end of the array. The latter version will often be faster as less branches occur.

In some cases we replace a pointer as in

```
class foo {
    ulong *x;
    foo(ulong n)  { x = new ulong[n]; /*...*/ }
```

by an array

```
class foo {
    ulong x[32];  // works for n<=32 only
    foo(ulong n)  { /*...*/ }
```

The latter version can be faster as the memory address can often be generated more cheaply. This optimization applies to most algorithms for the generation of permutations.

The test for equality as in

```
ulong k = 0;
while ( k!=n )  { /*...*/  ++k; }
```

is more expensive than the test for equality with zero as in

```
ulong k = n;
while ( --k!=0 )  { /*...*/ }
```

Therefore the latter version should be used when applicable.

To reduce the number of branches, replace two tests

```
if ( (x<0) || (x>m) )  { /*...*/ }
```

by the following single test where unsigned integers are used:

```
if ( x>m )  { /*...*/ }
```

Use a do-while construct instead of a while-do loop whenever possible because the latter also tests the loop condition at entry. Even when the do-while version causes some unnecessary work the gain from the avoided branch may outweigh it. Note that in the C language the for-loop also tests the condition at loop entry.

When computing the next object there may be special cases where the update is trivial. If the percentage of these 'easy cases' is not too small an extra branch in the update routine should be created. The performance gain is very visible in most cases (section 10.4 on page 245) and can be drastic (section 10.5 on page 248).

Recursive routines can very be elegant and versatile. However, expect only about half the speed of a good iterative implementation of the same algorithm. The ideas given in section 4.8 on page 164 can be used to obtain iterative versions from a recursive implementation.

Approximate timings for the routines are given with most implementations. These were obtained with the referenced demo-programs, simply uncomment the line

```
//#define TIMING // uncomment to disable printing
```

near the top of the program file for your own measurements. The parameters (program arguments) used for the timing are given near the end of the file. When a generator routine is used in an application one must do the benchmarking with the application. Further optimization will very likely involve the surrounding code rather than the generator alone. The rate of generation for a certain object is occasionally given as 123 M/s, meaning that 123 million objects are generated per second.

Choosing the optimal ordering and type of representation (for example, delta sets versus sets) for the given task is crucial.

Address generation can be simpler if arrays are used instead of pointers. The speedup gained can be spectacular, see section 7.1 on page 195.

5.5 Remarks about the C++ implementations

The iterative methods are implemented as C++ classes. The first object in the given order is created by the method `first()`. The method to compute the successor is usually `next()`. If a method for the computation of the predecessor is given then it is called `prev()`, and the method `last()` computes the last element in the list.

The combinatorial object can be accessed through the method `data()`. In order to make all data of a class accessible the data is declared `public`. Thereby the need for various `get_something()` methods is avoided. To minimize the danger of accidental modification of class data the variable names end with an underscore. For example, the class for the generation of combinations in lexicographic order starts as

```
class combination_lex
{
public:
    ulong *x_;    // combination: k elements 0<=x[j]<k in increasing order
    ulong n_, k_; // Combination (n choose k)
```

The methods for the user of the class are `public`, the internal methods are declared `private`.

Chapter 6

Combinations

n \ k	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0:	1															
1:	1	1														
2:	1	2	1													
3:	1	3	3	1												
4:	1	4	6	4	1											
5:	1	5	10	10	5	1										
6:	1	6	15	20	15	6	1									
7:	1	7	21	35	35	21	7	1								
8:	1	8	28	56	70	56	28	8	1							
9:	1	9	36	84	126	126	84	36	9	1						
10:	1	10	45	120	210	252	210	120	45	10	1					
11:	1	11	55	165	330	462	462	330	165	55	11	1				
12:	1	12	66	220	495	792	924	792	495	220	66	12	1			
13:	1	13	78	286	715	1287	1716	1716	1287	715	286	78	13	1		
14:	1	14	91	364	1001	2002	3003	3432	3003	2002	1001	364	91	14	1	
15:	1	15	105	455	1365	3003	5005	6435	6435	5005	3003	1365	455	105	15	1

Figure 6.0-A: The binomial coefficients $\binom{n}{k}$ for $0 \le n, k \le 15$.

The number of ways to choose k elements from a set of n elements equals the binomial coefficient ('n choose k'):

$$\binom{n}{k} \;=\; \frac{n!}{k!\,(n-k)!} \;=\; \frac{\prod_{j=1}^{k}(n-j+1)}{k!} \tag{6.0-1}$$

Equivalently, a set of n elements has $\binom{n}{k}$ subsets of exactly k elements. These subsets are called the *k-subsets* (where k is fixed) or *k-combinations* of an n-set.

To avoid overflow during the computation of the binomial coefficient, use the form

$$\binom{n}{k} \;=\; \frac{(n-k+1)^{\overline{k}}}{1^{\overline{k}}} \;=\; \frac{n-k+1}{1}\cdot\frac{n-k+2}{2}\cdot\frac{n-k+3}{3}\cdots\frac{n}{k} \tag{6.0-2}$$

An implementation is given in [FXT: aux0/binomial.h]:

```
1   inline ulong binomial(ulong n, ulong k)
2   {
3       if ( k>n )  return  0;
4       if ( (k==0) || (k==n) )  return  1;
5       if ( 2*k > n )  k = n-k;  // use symmetry
6
7       ulong b = n - k + 1;
8       ulong f = b;
9       for (ulong j=2; j<=k; ++j)
10      {
11          ++f;
12          b *= f;
```

```
13          b /= j;
14      }
15      return  b;
16  }
```

The table of the first binomial coefficients is shown in figure 6.0-A. This table is called *Pascal's triangle*, it was generated with the program [FXT: comb/binomial-demo.cc]. Observe that

$$\binom{n}{k} = \binom{n-1}{k-1} + \binom{n-1}{k} \tag{6.0-3}$$

That is, each entry is the sum of its upper and left upper neighbor.

6.1 Lexicographic and co-lexicographic order

lexicographic			co-lexicographic			
	set	delta set		set	delta set	set reversed
1:	{ 0, 1, 2 }	111...	1:	{ 0, 1, 2 }	111...	{ 2, 1, 0 }
2:	{ 0, 1, 3 }	11.1..	2:	{ 0, 1, 3 }	11.1..	{ 3, 1, 0 }
3:	{ 0, 1, 4 }	11..1.	3:	{ 0, 2, 3 }	1.11..	{ 3, 2, 0 }
4:	{ 0, 1, 5 }	11...1	4:	{ 1, 2, 3 }	.111..	{ 3, 2, 1 }
5:	{ 0, 2, 3 }	1.11..	5:	{ 0, 1, 4 }	11..1.	{ 4, 1, 0 }
6:	{ 0, 2, 4 }	1.1.1.	6:	{ 0, 2, 4 }	1.1.1.	{ 4, 2, 0 }
7:	{ 0, 2, 5 }	1.1..1	7:	{ 1, 2, 4 }	.11.1.	{ 4, 2, 1 }
8:	{ 0, 3, 4 }	1..11.	8:	{ 0, 3, 4 }	1..11.	{ 4, 3, 0 }
9:	{ 0, 3, 5 }	1..1.1	9:	{ 1, 3, 4 }	.1.11.	{ 4, 3, 1 }
10:	{ 0, 4, 5 }	1...11	10:	{ 2, 3, 4 }	..111.	{ 4, 3, 2 }
11:	{ 1, 2, 3 }	.111..	11:	{ 0, 1, 5 }	11...1	{ 5, 1, 0 }
12:	{ 1, 2, 4 }	.11.1.	12:	{ 0, 2, 5 }	1.1..1	{ 5, 2, 0 }
13:	{ 1, 2, 5 }	.11..1	13:	{ 1, 2, 5 }	.11..1	{ 5, 2, 1 }
14:	{ 1, 3, 4 }	.1.11.	14:	{ 0, 3, 5 }	1..1.1	{ 5, 3, 0 }
15:	{ 1, 3, 5 }	.1.1.1	15:	{ 1, 3, 5 }	.1.1.1	{ 5, 3, 1 }
16:	{ 1, 4, 5 }	.1..11	16:	{ 2, 3, 5 }	..11.1	{ 5, 3, 2 }
17:	{ 2, 3, 4 }	..111.	17:	{ 0, 4, 5 }	1...11	{ 5, 4, 0 }
18:	{ 2, 3, 5 }	..11.1	18:	{ 1, 4, 5 }	.1..11	{ 5, 4, 1 }
19:	{ 2, 4, 5 }	..1.11	19:	{ 2, 4, 5 }	..1.11	{ 5, 4, 2 }
20:	{ 3, 4, 5 }	...111	20:	{ 3, 4, 5 }	...111	{ 5, 4, 3 }

Figure 6.1-A: All combinations $\binom{6}{3}$ in lexicographic order (left), and co-lexicographic order (right).

The combinations of three elements out of six in *lexicographic* (or simply *lex*) order are shown in figure 6.1-A (left). The sequence is such that the sets are ordered lexicographically. Note that for the delta sets the element zero is printed first whereas with binary words (section 1.26 on page 66) the least significant bit (bit zero) is printed last. The sequence for *co-lexicographic* (or *colex*) order is such that the sets, when written reversed, are ordered lexicographically.

6.1.1 Lexicographic order

The C++ class [FXT: `class combination_lex` in comb/combination-lex.h] generates the sets:

```
1   class combination_lex
2   {
3   public:
4       ulong *x_;    // combination: k elements 0<=x[j]<k in increasing order
5       ulong n_, k_; // Combination (n choose k)
6
7   public:
8       combination_lex(ulong n, ulong k)
9       {
10          n_ = n;  k_ = k;
11          x_ = new ulong[k_];
```

```
12          first();
13      }
14
15      ~combination_lex()  { delete [] x_; }
16
17      void first()
18      {
19          for (ulong k=0; k<k_; ++k)  x_[k] = k;
20      }
21
22      void last()
23      {
24          for (ulong i=0; i<k_; ++i)  x_[i] = n_ - k_ + i;
25      }
26
```

Computation of the successor and predecessor:

```
1       ulong next()
2       // Return smallest position that changed, return k with last combination
3       {
4           if ( x_[0] == n_ - k_ )  // current combination is the last
5           { first();  return k_; }
6
7           ulong j = k_ - 1;
8           // easy case:  highest element != highest possible value:
9           if ( x_[j] < (n_-1) )  { ++x_[j];  return j; }
10
11          // find highest falling edge:
12          while ( 1 == (x_[j] - x_[j-1]) )  { --j; }
13
14          // move lowest element of highest block up:
15          ulong ret = j - 1;
16          ulong z = ++x_[j-1];
17
18          // ... and attach rest of block:
19          while ( j < k_ )  { x_[j] = ++z;  ++j; }
20
21          return  ret;
22      }
```

```
1       ulong prev()
2       // Return smallest position that changed, return k with last combination
3       {
4           if ( x_[k_-1] == k_-1 )  // current combination is the first
5           { last();  return k_; }
6
7           // find highest falling edge:
8           ulong j = k_ - 1;
9           while ( 1 == (x_[j] - x_[j-1]) )  { --j; }
10
11          ulong ret = j;
12          --x_[j];  // move down edge element
13
14          // ... and move rest of block to high end:
15          while ( ++j < k_ )  x_[j] = n_ - k_ + j;
16
17          return  ret;
18      }
```

The listing in figure 6.1-A was created with the program [FXT: comb/combination-lex-demo.cc]. The routine generates the combinations $\binom{32}{20}$ at a rate of about 95 million per second. The combinations $\binom{32}{12}$ are generated at a rate of 160 million per second.

6.1.2 Co-lexicographic order

The combinations of three elements out of six in *co-lexicographic* (or *colex*) order are shown in figure 6.1-A (right). Algorithms to compute the successor and predecessor are implemented in [FXT: `class combination_colex` in comb/combination-colex.h]:

```
1   class combination_colex
2   {
3   public:
4       ulong *x_;    // combination: k elements 0<=x[j]<k in increasing order
```

```
5       ulong n_, k_; // Combination (n choose k)
6
7       combination_colex(ulong n, ulong k)
8       {
9           n_ = n;  k_ = k;
10          x_ = new ulong[k_+1];
11          x_[k_] = n_ + 2;  // sentinel
12          first();
13      }
14
15     [--snip--]
16      ulong next()
17      // Return greatest position that changed, return k with last combination
18      {
19
20          if ( x_[0] == n_ - k_ )  // current combination is the last
21          { first();  return k_; }
22
23          ulong j = 0;
24          // until lowest rising edge:  attach block at low end
25          while ( 1 == (x_[j+1] - x_[j]) )  { x_[j] = j;  ++j; }  // can touch sentinel
26
27          ++x_[j];  // move edge element up
28
29          return  j;
30      }
31
32      ulong prev()
33      // Return greatest position that changed, return k with last combination
34      {
35          if ( x_[k_-1] == k_-1 )  // current combination is the first
36          { last();  return k_; }
37
38          // find lowest falling edge:
39          ulong j = 0;
40          while ( j == x_[j] )  ++j;  // can touch sentinel
41
42          --x_[j];  // move edge element down
43          ulong ret = j;
44
45          // attach rest of low block:
46          while ( 0!=j-- )  x_[j] = x_[j+1] - 1;
47
48          return  ret;
49      }
50     [--snip--]
```

The listing in figure 6.1-A was created with the program [FXT: comb/combination-colex-demo.cc]. The combinations are generated $\binom{32}{20}$ at a rate of about 140 million objects per second, the combinations $\binom{32}{12}$ are generated at a rate of 190 million objects per second.

As a toy application of the combinations in co-lexicographic order we compute the products of k of the n smallest primes. We maintain an array of k products shown at the right of figure 6.1-B. When the return value of the method `next()` is j then $j+1$ elements have to be updated from right to left [FXT: comb/kproducts-colex-demo.cc]:

```
1       combination_colex C(n, k);
2       const ulong *c = C.data();  // combinations as sets
3
4       ulong *tf = new ulong[n];  // table of Factors (primes)
5       // fill in small primes:
6       for (ulong j=0,f=2; j<n; ++j)  { tf[j] = f;  f=next_small_prime(f+1); }
7
8       ulong *tp = new ulong[k+1];  // table of Products
9       tp[k] = 1; // one appended (sentinel)
10
11      ulong j = k-1;
12      do
13      {
14          // update products from right:
15          ulong x = tp[j+1];
16          { ulong i = j;
17          do
18          {
19              ulong f = tf[ c[i] ];
20              x *= f;
```

```
              combination  j     delta-set      products
        1:    { 0, 1, 2 }  2      111....    [   30   15    5    1 ]
        2:    { 0, 1, 3 }  2      11.1...    [   42   21    7    1 ]
        3:    { 0, 2, 3 }  1      1.11...    [   70   35    7    1 ]
        4:    { 1, 2, 3 }  0      .111...    [  105   35    7    1 ]
        5:    { 0, 1, 4 }  2      11..1..    [   66   33   11    1 ]
        6:    { 0, 2, 4 }  1      1.1.1..    [  110   55   11    1 ]
        7:    { 1, 2, 4 }  0      .11.1..    [  165   55   11    1 ]
        8:    { 0, 3, 4 }  1      1..11..    [  154   77   11    1 ]
        9:    { 1, 3, 4 }  0      .1.11..    [  231   77   11    1 ]
       10:    { 2, 3, 4 }  0      ..111..    [  385   77   11    1 ]
       11:    { 0, 1, 5 }  2      11...1.    [   78   39   13    1 ]
       12:    { 0, 2, 5 }  1      1.1..1.    [  130   65   13    1 ]
       13:    { 1, 2, 5 }  0      .11..1.    [  195   65   13    1 ]
       14:    { 0, 3, 5 }  1      1..1.1.    [  182   91   13    1 ]
       15:    { 1, 3, 5 }  0      .1.1.1.    [  273   91   13    1 ]
       16:    { 2, 3, 5 }  0      ..11.1.    [  455   91   13    1 ]
       17:    { 0, 4, 5 }  1      1...11.    [  286  143   13    1 ]
       18:    { 1, 4, 5 }  0      .1..11.    [  429  143   13    1 ]
       19:    { 2, 4, 5 }  0      ..1.11.    [  715  143   13    1 ]
       20:    { 3, 4, 5 }  0      ...111.    [ 1001  143   13    1 ]
       21:    { 0, 1, 6 }  2      11....1    [  102   51   17    1 ]
       22:    { 0, 2, 6 }  1      1.1...1    [  170   85   17    1 ]
       23:    { 1, 2, 6 }  0      .11...1    [  255   85   17    1 ]
       24:    { 0, 3, 6 }  1      1..1..1    [  238  119   17    1 ]
       25:    { 1, 3, 6 }  0      .1.1..1    [  357  119   17    1 ]
       26:    { 2, 3, 6 }  0      ..11..1    [  595  119   17    1 ]
       27:    { 0, 4, 6 }  1      1...1.1    [  374  187   17    1 ]
       28:    { 1, 4, 6 }  0      .1..1.1    [  561  187   17    1 ]
       29:    { 2, 4, 6 }  0      ..1.1.1    [  935  187   17    1 ]
       30:    { 3, 4, 6 }  0      ...11.1    [ 1309  187   17    1 ]
       31:    { 0, 5, 6 }  1      1....11    [  442  221   17    1 ]
       32:    { 1, 5, 6 }  0      .1...11    [  663  221   17    1 ]
       33:    { 2, 5, 6 }  0      ..1..11    [ 1105  221   17    1 ]
       34:    { 3, 5, 6 }  0      ...1.11    [ 1547  221   17    1 ]
       35:    { 4, 5, 6 }  0      ....111    [ 2431  221   17    1 ]
```

Figure 6.1-B: All products of $k = 3$ of the $n = 7$ smallest primes $(2, 3, 5, \ldots, 17)$. The products are the leftmost elements of the array on the right hand side.

```
21              tp[i] = x;
22          }
23          while ( i-- );
24          } // here: final product is x == tp[0]
25
26          // visit the product x here
27
28          j = C.next();
29      }
30      while ( j < k );
```

The leftmost element of this array is the desired product. A sentinel element one at the end of the array is used to keep the code simple. With lexicographic order the update would go from left to right.

6.2 Order by prefix shifts (cool-lex)

An algorithm for generating combinations by prefix shifts which is given in [226]. The ordering is called *cool-lex* in the paper. Figure 6.2-A shows some orders for $\binom{5}{k}$, figure 6.2-B shows the combinations $\binom{9}{3}$. The listings were created with the program [FXT: comb/combination-pref-demo.cc] which uses the implementation in [FXT: `class combination_pref` in comb/combination-pref.h]:

```
1   class combination_pref
2   {
3   public:
```

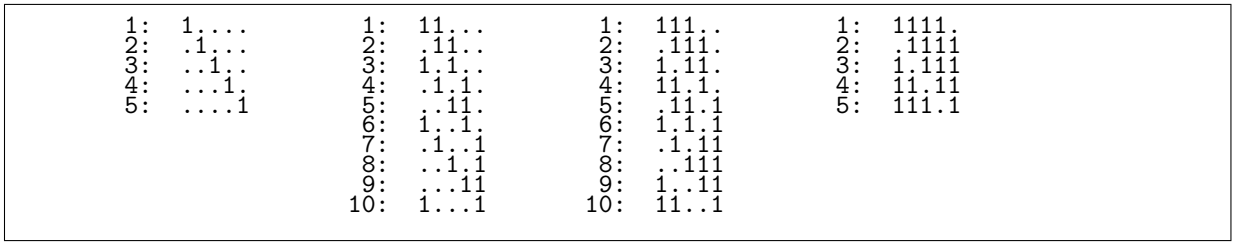

Figure 6.2-A: Combinations $\binom{5}{k}$, for $k = 1, 2, 3, 4$ in an ordering generated by prefix shifts.

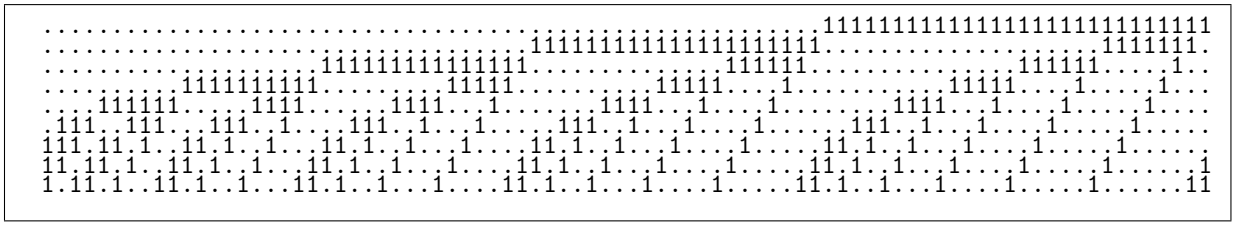

Figure 6.2-B: Combinations $\binom{9}{3}$ via prefix shifts.

```
4       ulong *b_;  // data as delta set
5       ulong s_, t_, n_;  // combination (n choose k) where n=s+t, k=t.
6   private:
7       ulong x, y;  // aux
8
9   public:
10      combination_pref(ulong n, ulong k)
11      // Must have: n>=2, k>=1  (i.e. s!=0 and t!=0)
12      {
13          s_ = n - k;
14          t_ = k;
15          n_ = s_ + t_;
16          b_ = new ulong[n_];
17          first();
18      }
19    [--snip--]
```

```
1       void first()
2       {
3           for (ulong j=0; j<n_; ++j)  b_[j] = 0;
4           for (ulong j=0; j<t_; ++j)  b_[j] = 1;
5           x = 0;  y = 0;
6       }
7
8       bool next()
9       {
10          if ( x==0 )  { x=1;  b_[t_]=1;  b_[0]=0;  return true; }
11          else
12          {
13              if ( x>=n_-1 )  return false;
14              else
15              {
16                  b_[x] = 0; ++x;  b_[y] = 1; ++y; // X(s,t)
17                  if ( b_[x]==0 )
18                  {
19                      b_[x] = 1;  b_[0] = 0;       // Y(s,t)
20                      if ( y>1 )  x = 1;           // Z(s,t)
21                      y = 0;
22                  }
23                  return true;
24              }
25          }
26      }
27    [--snip--]
```

The combinations are generated $\binom{32}{20}$ at a rate of about 95 million objects per second, the combinations $\binom{32}{12}$ are generated at a rate of 85 million objects per second.

6.3 Minimal-change order

Gray code			complemented Gray code		
1:	{ 0, 1, 2 }	111...	1:	{ 3, 4, 5 }	...111
2:	{ 0, 2, 3 }	1.11..	2:	{ 1, 4, 5 }	.1..11
3:	{ 1, 2, 3 }	.111..	3:	{ 0, 4, 5 }	1...11
4:	{ 0, 1, 3 }	11.1..	4:	{ 2, 4, 5 }	..1.11
5:	{ 0, 3, 4 }	1..11.	5:	{ 1, 2, 5 }	.11..1
6:	{ 1, 3, 4 }	.1.11.	6:	{ 0, 2, 5 }	1.1..1
7:	{ 2, 3, 4 }	..111.	7:	{ 0, 1, 5 }	11...1
8:	{ 0, 2, 4 }	1.1.1.	8:	{ 1, 3, 5 }	.1.1.1
9:	{ 1, 2, 4 }	.11.1.	9:	{ 0, 3, 5 }	1..1.1
10:	{ 0, 1, 4 }	11..1.	10:	{ 2, 3, 5 }	..11.1
11:	{ 0, 4, 5 }	1...11	11:	{ 1, 2, 3 }	.111..
12:	{ 1, 4, 5 }	.1..11	12:	{ 0, 2, 3 }	1.11..
13:	{ 2, 4, 5 }	..1.11	13:	{ 0, 1, 3 }	11.1..
14:	{ 3, 4, 5 }	...111	14:	{ 0, 1, 2 }	111...
15:	{ 0, 3, 5 }	1..1.1	15:	{ 1, 2, 4 }	.11.1.
16:	{ 1, 3, 5 }	.1.1.1	16:	{ 0, 2, 4 }	1.1.1.
17:	{ 2, 3, 5 }	..11.1	17:	{ 0, 1, 4 }	11..1.
18:	{ 0, 2, 5 }	1.1..1	18:	{ 1, 3, 4 }	.1.11.
19:	{ 1, 2, 5 }	.11..1	19:	{ 0, 3, 4 }	1..11.
20:	{ 0, 1, 5 }	11...1	20:	{ 2, 3, 4 }	..111.

Figure 6.3-A: Combinations $\binom{6}{3}$ in Gray order (left), and complemented Gray order (right).

The combinations of three elements out of six in a *minimal-change* order (a *Gray code*) are shown in figure 6.3-A (left). With each transition exactly one element changes its position. We use a recursion for the list $C(n,k)$ of combinations $\binom{n}{k}$ (symbols as in relation 12.0-1 on page 295):

$$C(n,k) \;=\; \begin{matrix}[C(n-1,\,k) \qquad\quad\;\;] \\ [(n)\,.\,C^{\mathbf{R}}(n-1,\,k-1)]\end{matrix} \;=\; \begin{matrix}[0\,.\,C(n-1,\,k) \qquad\;\;] \\ [1\,.\,C^{\mathbf{R}}(n-1,\,k-1)]\end{matrix} \tag{6.3-1}$$

The first equality is for the set representation, the second for the delta-set representation. An implementation is given in [FXT: comb/combination-gray-rec-demo.cc]:

```
1   ulong *x;   // elements in combination at x[1] ... x[k]
2
3   void comb_gray(ulong n, ulong k, bool z)
4   {
5       if ( k==n )
6       {
7           for (ulong j=1; j<=k; ++j)  x[j] = j;
8           visit();
9           return;
10      }
11
12      if ( z )  // forward:
13      {
14          comb_gray(n-1, k, z);
15          if ( k>0 )  { x[k] = n;  comb_gray(n-1, k-1, !z); }
16      }
17      else     // backward:
18      {
19          if ( k>0 )  { x[k] = n;  comb_gray(n-1, k-1, !z); }
20          comb_gray(n-1, k, z);
21      }
22  }
```

A recursion for the complemented order is

$$C(n,k) \;=\; \begin{matrix}[(n)\,.\,C(n-1,\,k-1)] \\ [C^{\mathbf{R}}(n-1,\,k) \qquad\;\;]\end{matrix} \;=\; \begin{matrix}[1\,.\,C(n-1,\,k-1)] \\ [0\,.\,C^{\mathbf{R}}(n-1,\,k) \;\;]\end{matrix} \tag{6.3-2}$$

```
1   void comb_gray_compl(ulong n, ulong k, bool z)
2   {
3     [--snip--]
```

```
4       if ( z )  // forward:
5       {
6           if ( k>0 )  { x[k] = n;  comb_gray_compl(n-1, k-1, z); }
7           comb_gray_compl(n-1, k, !z);
8       }
9       else     // backward:
10      {
11          comb_gray_compl(n-1, k, !z);
12          if ( k>0 )  { x[k] = n;  comb_gray_compl(n-1, k-1, z); }
13      }
14  }
```

A very efficient (*revolving door*) algorithm that generates the sets for the Gray code is given in [212]. An implementation following [170] is [FXT: `class combination_revdoor` in comb/combination-revdoor.h]. Usage of the class is shown in [FXT: comb/combination-revdoor-demo.cc]. The routine generates the combinations $\binom{32}{20}$ at a rate of about 115 M/s, the combinations $\binom{32}{12}$ are generated at a rate of 181 M/s. An implementation geared for good performance with small values of k is given in [175], a C++ adaptation is [FXT: comb/combination-lam-demo.cc]. The combinations $\binom{32}{12}$ a generated at a rate of 190 M/s, and 250 M/s for the combinations $\binom{64}{7}$. The routine is limited to values $k \geq 2$.

6.4 The Eades-McKay strong minimal-change order

Eades-McKay		complemented Eades-McKay	
1:	{ 4, 5, 6 } 111	1:	{ 4, 5, 6 } 111
2:	{ 3, 5, 6 } ...1.11	2:	{ 3, 5, 6 } ...1.11
3:	{ 2, 5, 6 } ..1..11	3:	{ 2, 5, 6 } ..1..11
4:	{ 1, 5, 6 } .1...11	4:	{ 1, 5, 6 } .1...11
5:	{ 0, 5, 6 } 1....11	5:	{ 0, 5, 6 } 1....11
6:	{ 0, 1, 6 } 11....1	6:	{ 0, 4, 6 } 1...1.1
7:	{ 0, 2, 6 } 1.1...1	7:	{ 1, 4, 6 } .1..1.1
8:	{ 1, 2, 6 } .11...1	8:	{ 2, 4, 6 } ..1.1.1
9:	{ 1, 3, 6 } .1.1..1	9:	{ 3, 4, 6 } ...11.1
10:	{ 0, 3, 6 } 1..1..1	10:	{ 2, 3, 6 } ..11..1
11:	{ 2, 3, 6 } ..11..1	11:	{ 1, 3, 6 } .1.1..1
12:	{ 2, 4, 6 } ..1.1.1	12:	{ 0, 3, 6 } 1..1..1
13:	{ 1, 4, 6 } .1..1.1	13:	{ 0, 2, 6 } 1.1...1
14:	{ 0, 4, 6 } 1...1.1	14:	{ 1, 2, 6 } .11...1
15:	{ 3, 4, 6 } ...11.1	15:	{ 0, 1, 6 } 11....1
16:	{ 3, 4, 5 } ...111.	16:	{ 0, 1, 5 } 11...1.
17:	{ 2, 4, 5 } ..1.11.	17:	{ 0, 2, 5 } 1.1..1.
18:	{ 1, 4, 5 } .1..11.	18:	{ 1, 2, 5 } .11..1.
19:	{ 0, 4, 5 } 1...11.	19:	{ 2, 3, 5 } ..11.1.
20:	{ 0, 1, 5 } 11...1.	20:	{ 1, 3, 5 } .1.1.1.
21:	{ 0, 2, 5 } 1.1..1.	21:	{ 0, 3, 5 } 1..1.1.
22:	{ 1, 2, 5 } .11..1.	22:	{ 0, 4, 5 } 1...11.
23:	{ 1, 3, 5 } .1.1.1.	23:	{ 1, 4, 5 } .1..11.
24:	{ 0, 3, 5 } 1..1.1.	24:	{ 2, 4, 5 } ..1.11.
25:	{ 2, 3, 5 } ..11.1.	25:	{ 3, 4, 5 } ...111.
26:	{ 2, 3, 4 } ..111..	26:	{ 2, 3, 4 } ..111..
27:	{ 1, 3, 4 } .1.11..	27:	{ 1, 3, 4 } .1.11..
28:	{ 0, 3, 4 } 1..11..	28:	{ 0, 3, 4 } 1..11..
29:	{ 0, 1, 4 } 11..1..	29:	{ 0, 2, 4 } 1.1.1..
30:	{ 0, 2, 4 } 1.1.1..	30:	{ 1, 2, 4 } .11.1..
31:	{ 1, 2, 4 } .11.1..	31:	{ 0, 1, 4 } 11..1..
32:	{ 1, 2, 3 } .111...	32:	{ 0, 1, 3 } 11.1...
33:	{ 0, 2, 3 } 1.11...	33:	{ 0, 2, 3 } 1.11...
34:	{ 0, 1, 3 } 11.1...	34:	{ 1, 2, 3 } .111...
35:	{ 0, 1, 2 } 111....	35:	{ 0, 1, 2 } 111....

Figure 6.4-A: Combinations $\binom{6}{3}$ in the Eades-McKay order (left), and in the complemented Eades-Mckay order (right).

In any Gray code order for combinations just one element is moved between to successive combinations.

When an element is moved across any other then there is more than one change on the set representation. If i elements are crossed then $i+1$ entries in the set change:

```
        set          delta set
{ 0, 1, 2, 3 }    1111..
{ 1, 2, 3, 4 }    .1111.
```

A *strong minimal-change order* is a Gray code where only one entry in the set representation is changed per step. That is, only zeros in the delta set representation are crossed, and the moves are called *homogeneous*. One such order is the *Eades-McKay* sequence described in [110]. The Eades-McKay sequence for the combinations $\binom{6}{3}$ is shown in figure 6.4-A (left). It can be generated with the program [FXT: comb/combination-emk-rec-demo.cc]:

```
1   ulong *rv;  // elements in combination at rv[1] ... rv[k]
2
3   void
4   comb_emk(ulong n, ulong k, bool z)
5   {
6       if ( k==n )
7       {
8           for (ulong j=1; j<=k; ++j)  rv[j] = j;
9           visit();
10          return;
11      }
12
13      if ( z )  // forward:
14      {
15          if ( (n>=2) && (k>=2) )  { rv[k] = n;  rv[k-1] = n-1;  comb_emk(n-2, k-2, z); }
16          if ( (n>=2) && (k>=1) )  { rv[k] = n;  comb_emk(n-2, k-1, !z); }
17          if ( (n>=1) )            { comb_emk(n-1, k, z); }
18      }
19      else     // backward:
20      {
21          if ( (n>=1) )            { comb_emk(n-1, k, z); }
22          if ( (n>=2) && (k>=1) )  { rv[k] = n;  comb_emk(n-2, k-1, !z); }
23          if ( (n>=2) && (k>=2) )  { rv[k] = n;  rv[k-1] = n-1;  comb_emk(n-2, k-2, z); }
24      }
25  }
```

The combinations $\binom{32}{20}$ are generated at a rate of about 44 million per second, the combinations $\binom{32}{12}$ at a rate of 34 million per second.

The underlying recursion for the list $E(n,k)$ of combinations $\binom{n}{k}$ is (notations as in relation 12.0-1 on page 295)

$$E(n,k) \;=\; \begin{matrix}[(n)\,.\,(n-1)\,.\,E(n-2,\,k-2)] \\ [(n)\,.\,E^{\mathbf{R}}(n-2,\,k-1) \quad] \\ [E(n-1,\,k) \quad]\end{matrix} \;=\; \begin{matrix}[1\,1\,.\,E(n-2,\,k-2) \;] \\ [1\,0\,.\,E^{\mathbf{R}}(n-2,\,k-1)] \\ [0\,.\,E(n-1,\,k) \quad]\end{matrix} \tag{6.4-1}$$

Again, the first equality is for the set representation, the second for the delta-set representation. Counting the elements on both sides gives the relation

$$\binom{n}{k} \;=\; \binom{n-2}{k-2} + \binom{n-2}{k-1} + \binom{n-1}{k} \tag{6.4-2}$$

which is an easy consequence of relation 6.0-3 on page 176. A recursion for the complemented sequence (with respect to the delta sets) is

$$E'(n,k) \;=\; \begin{matrix}[E'(n-1,\,k-1) \quad] \\ [(n-1)\,.\,E'^{\mathbf{R}}(n-2,\,k-1)] \\ [(n)\,.\;.\,E'(n-2,\,k) \quad]\end{matrix} \;=\; \begin{matrix}[1\,.\,E'(n-1,\,k-1) \;] \\ [0\,1\,.\,E'^{\mathbf{R}}(n-2,\,k-1)] \\ [0\,0\,.\,E'(n-2,\,k) \quad]\end{matrix} \tag{6.4-3}$$

Counting on both sides gives

$$\binom{n}{k} \;=\; \binom{n-2}{k} + \binom{n-2}{k-1} + \binom{n-1}{k-1} \tag{6.4-4}$$

The condition for the recursion end has to be modified:

```
void
comb_emk_compl(ulong n, ulong k, bool z)
{
    if ( (k==0) || (k==n) )
    {
        for (ulong j=1; j<=k; ++j)  rv[j] = j;
        ++ct;
        visit();
        return;
    }

    if ( z )  // forward:
    {
        if ( (n>=1) && (k>=1) )  { rv[k] = n;  comb_emk_compl(n-1, k-1, z); }     // 1
        if ( (n>=2) && (k>=1) )  { rv[k] = n-1;  comb_emk_compl(n-2, k-1, !z); }  // 01
        if ( (n>=2)  )           { comb_emk_compl(n-2, k-0, z); }                 // 00
    }
    else     // backward:
    {
        if ( (n>=2) )            { comb_emk_compl(n-2, k-0, z); }                 // 00
        if ( (n>=2) && (k>=1) )  { rv[k] = n-1;  comb_emk_compl(n-2, k-1, !z); }  // 01
        if ( (n>=1) && (k>=1) )  { rv[k] = n;  comb_emk_compl(n-1, k-1, z); }     // 1
    }
}
```

The complemented sequence is not a strong Gray code.

Iterative generation via modulo moves

An iterative algorithm for the Eades-McKay sequence is given in [FXT: `class combination_emk` in comb/combination-emk.h]:

```
class combination_emk
{
public:
    ulong *x_;  // combination: k elements 0<=x[j]<k in increasing order
    ulong *s_;  // aux: start of range for moves
    ulong *a_;  // aux: actual start position of moves
    ulong n_, k_; // Combination (n choose k)

public:
    combination_emk(ulong n, ulong k)
    {
        n_ = n;
        k_ = k;
        x_ = new ulong[k_+1];  // incl. high sentinel
        s_ = new ulong[k_+1];  // incl. high sentinel
        a_ = new ulong[k_];
        x_[k_] = n_;
        first();
    }
  [--snip--]

    void first()
    {
        for (ulong j=0; j<k_; ++j)  x_[j] = j;
        for (ulong j=0; j<k_; ++j)  s_[j] = j;
        for (ulong j=0; j<k_; ++j)  a_[j] = x_[j];
    }
```

The computation of the successor uses modulo steps:

```
    ulong next()
    // Return position where track changed, return k with last combination
    {
        ulong j = k_;
        while ( j-- )  // loop over tracks
        {
            const ulong sj = s_[j];
            const ulong m = x_[j+1] - sj - 1;

            if ( 0!=m )  // unless range empty
            {
```

```
12                  ulong u = x_[j] - sj;
13
14                  // modulo moves:
15                  if ( 0==(j&1) )
16                  {
17                      ++u;
18                      if ( u>m )  u = 0;
19                  }
20                  else
21                  {
22                      --u;
23                      if ( u>m )  u = m;
24                  }
25                  u += sj;
26
27                  if ( u != a_[j] )  // next pos != start position
28                  {
29                      x_[j] = u;
30                      s_[j+1] = u+1;
31                      return j;
32                  }
33              }
34              a_[j] = x_[j];
35          }
36
37          return  k_;  // current combination is last
38      }
39  };
```

The combinations $\binom{32}{20}$ are generated at a rate of about 60 million per second, the combinations $\binom{32}{12}$ at a rate of 85 million per second [FXT: comb/combination-emk-demo.cc].

6.5 Two-close orderings via endo/enup moves

6.5.1 The endo and enup orderings for numbers

m	endo sequence	m	enup sequence
1:	1 0	1:	0 1
2:	1 2 0	2:	0 2 1
3:	1 3 2 0	3:	0 2 3 1
4:	1 3 4 2 0	4:	0 2 4 3 1
5:	1 3 5 4 2 0	5:	0 2 4 5 3 1
6:	1 3 5 6 4 2 0	6:	0 2 4 6 5 3 1
7:	1 3 5 7 6 4 2 0	7:	0 2 4 6 7 5 3 1
8:	1 3 5 7 8 6 4 2 0	8:	0 2 4 6 8 7 5 3 1
9:	1 3 5 7 9 8 6 4 2 0	9:	0 2 4 6 8 9 7 5 3 1

Figure 6.5-A: The endo (left), and enup (right) orderings with maximal value m.

The *endo* order of the set $\{0, 1, 2, \ldots, m\}$ is obtained by writing all odd numbers of the set in increasing order followed by all even numbers in decreasing order: $\{1, 3, 5, \ldots, 6, 4, 2, 0\}$. The term endo stands for 'Even Numbers DOwn, odd numbers up'. A routine for generating the successor in endo order with maximal value m is [FXT: comb/endo-enup.h]:

```
1   inline ulong next_endo(ulong x, ulong m)
2   // Return next number in endo order
3   {
4       if ( x & 1 )  // x odd
5       {
6           x += 2;
7           if ( x>m )  x = m - (m&1);  // == max even <= m
8       }
9       else  // x even
10      {
11          x = ( x==0 ? 1 : x-2 );
12      }
13      return x;
14  }
```

The sequences for the first few m are shown in figure 6.5-A. The routine computes one for the input zero.

An ordering starting with the even numbers in increasing order will be called *enup* (for 'Even Numbers UP, odd numbers down'). The computation of the successor can be implemented as

```
1   static inline ulong next_enup(ulong x, ulong m)
2   {
3       if ( x & 1 )  // x odd
4       {
5           x = ( x==1 ? 0 : x-2 );
6       }
7       else  // x even
8       {
9           x += 2;
10          if ( x>m )  x = m - !(m&1);  // max odd <=m
11      }
12      return x;
13  }
```

As the orderings are reversals of each other, we define:

```
1   static inline ulong prev_endo(ulong x, ulong m)  { return next_enup(x, m); }
2   static inline ulong prev_enup(ulong x, ulong m)  { return next_endo(x, m); }
```

A function that returns the x-th number in enup order with maximal digit m is

```
1   static inline ulong enup_num(ulong x, ulong m)
2   {
3       ulong r = 2*x;
4       if  ( r>m )  r = 2*m+1 - r;
5       return r;
6   }
```

The function will only work if $x \leq m$. For example, with $m = 5$:

```
x:   0 1 2 3 4 5
r:   0 2 4 5 3 1
```

The inverse function is

```
1   static inline ulong enup_idx(ulong x, ulong m)
2   {
3       const ulong b = x & 1;
4       x >>= 1;
5       return  ( b ? m-x : x );
6   }
```

The equivalent function to map into endo order is

```
1   static inline ulong endo_num(ulong x, ulong m)
2   {
3       // return enup_num(m-x, m);
4       x = m - x;
5       ulong r = 2*x;
6       if  ( r>m )  r = 2*m+1 - r;
7       return  r;
8   }
```

For example,

```
x:   0 1 2 3 4 5
r:   1 3 5 4 2 0
```

The inverse is

```
1   static inline ulong endo_idx(ulong x, ulong m)
2   {
3       const ulong b = x & 1;
4       x >>= 1;
5       return  ( b ? x : m-x );
6   }
```

6.5.2 The endo and enup orderings

Two strong minimal-change orderings for combinations can be obtained via moves in enup and endo order. Figure 6.5-B shows an ordering where the moves to the right are on even positions (enup order, left). If the moves to the right are on odd positions (endo order) then Chase's sequence is obtained (right). Both have the property of being *two-close*: an element in the delta set moves by at most two positions (and

	enup moves			endo moves	
1:	{ 0, 1, 2 }	111.....	1:	{ 0, 1, 2 }	111.....
2:	{ 0, 1, 4 }	11..1...	2:	{ 0, 1, 3 }	11.1....
3:	{ 0, 1, 6 }	11....1.	3:	{ 0, 1, 5 }	11...1..
4:	{ 0, 1, 7 }	11.....1	4:	{ 0, 1, 7 }	11.....1
5:	{ 0, 1, 5 }	11...1..	5:	{ 0, 1, 6 }	11....1.
6:	{ 0, 1, 3 }	11.1....	6:	{ 0, 1, 4 }	11..1...
7:	{ 0, 2, 3 }	1.11....	7:	{ 0, 3, 4 }	1..11...
8:	{ 0, 2, 4 }	1.1.1...	8:	{ 0, 3, 5 }	1..1.1..
9:	{ 0, 2, 6 }	1.1...1.	9:	{ 0, 3, 7 }	1..1...1
10:	{ 0, 2, 7 }	1.1....1	10:	{ 0, 3, 6 }	1..1..1.
11:	{ 0, 2, 5 }	1.1..1..	11:	{ 0, 5, 6 }	1....11.
12:	{ 0, 4, 5 }	1...11..	12:	{ 0, 5, 7 }	1....1.1
13:	{ 0, 4, 6 }	1...1.1.	13:	{ 0, 6, 7 }	1.....11
14:	{ 0, 4, 7 }	1...1..1	14:	{ 0, 4, 7 }	1...1..1
15:	{ 0, 6, 7 }	1.....11	15:	{ 0, 4, 6 }	1...1.1.
16:	{ 0, 5, 7 }	1....1.1	16:	{ 0, 4, 5 }	1...11..
17:	{ 0, 5, 6 }	1....11.	17:	{ 0, 2, 5 }	1.1..1..
18:	{ 0, 3, 6 }	1..1..1.	18:	{ 0, 2, 7 }	1.1....1
19:	{ 0, 3, 7 }	1..1...1	19:	{ 0, 2, 6 }	1.1...1.
20:	{ 0, 3, 5 }	1..1.1..	20:	{ 0, 2, 4 }	1.1.1...
21:	{ 0, 3, 4 }	1..11...	21:	{ 0, 2, 3 }	1.11....
22:	{ 2, 3, 4 }	..111...	22:	{ 1, 2, 3 }	.111....
23:	{ 2, 3, 6 }	..11..1.	23:	{ 1, 2, 5 }	.11..1..
24:	{ 2, 3, 7 }	..11...1	24:	{ 1, 2, 7 }	.11....1
25:	{ 2, 3, 5 }	..11.1..	25:	{ 1, 2, 6 }	.11...1.
26:	{ 2, 4, 5 }	..1.11..	26:	{ 1, 2, 4 }	.11.1...
27:	{ 2, 4, 6 }	..1.1.1.	27:	{ 1, 3, 4 }	.1.11...
28:	{ 2, 4, 7 }	..1.1..1	28:	{ 1, 3, 5 }	.1.1.1..
29:	{ 2, 6, 7 }	..1...11	29:	{ 1, 3, 7 }	.1.1...1
30:	{ 2, 5, 7 }	..1..1.1	30:	{ 1, 3, 6 }	.1.1..1.
31:	{ 2, 5, 6 }	..1..11.	31:	{ 1, 5, 6 }	.1...11.
32:	{ 4, 5, 6 }	....111.	32:	{ 1, 5, 7 }	.1...1.1
33:	{ 4, 5, 7 }	....11.1	33:	{ 1, 6, 7 }	.1....11
34:	{ 4, 6, 7 }	....1.11	34:	{ 1, 4, 7 }	.1..1..1
35:	{ 5, 6, 7 }	111	35:	{ 1, 4, 6 }	.1..1.1.
36:	{ 3, 6, 7 }	...1..11	36:	{ 1, 4, 5 }	.1..11..
37:	{ 3, 5, 7 }	...1.1.1	37:	{ 3, 4, 5 }	...111..
38:	{ 3, 5, 6 }	...1.11.	38:	{ 3, 4, 7 }	...11..1
39:	{ 3, 4, 6 }	...11.1.	39:	{ 3, 4, 6 }	...11.1.
40:	{ 3, 4, 7 }	...11..1	40:	{ 3, 5, 6 }	...1.11.
41:	{ 3, 4, 5 }	...111..	41:	{ 3, 5, 7 }	...1.1.1
42:	{ 1, 4, 5 }	.1..11..	42:	{ 3, 6, 7 }	...1..11
43:	{ 1, 4, 6 }	.1..1.1.	43:	{ 5, 6, 7 }	111
44:	{ 1, 4, 7 }	.1..1..1	44:	{ 4, 6, 7 }	....1.11
45:	{ 1, 6, 7 }	.1....11	45:	{ 4, 5, 7 }	....11.1
46:	{ 1, 5, 7 }	.1...1.1	46:	{ 4, 5, 6 }	....111.
47:	{ 1, 5, 6 }	.1...11.	47:	{ 2, 5, 6 }	..1..11.
48:	{ 1, 3, 6 }	.1.1..1.	48:	{ 2, 5, 7 }	..1..1.1
49:	{ 1, 3, 7 }	.1.1...1	49:	{ 2, 6, 7 }	..1...11
50:	{ 1, 3, 5 }	.1.1.1..	50:	{ 2, 4, 7 }	..1.1..1
51:	{ 1, 3, 4 }	.1.11...	51:	{ 2, 4, 6 }	..1.1.1.
52:	{ 1, 2, 4 }	.11.1...	52:	{ 2, 4, 5 }	..1.11..
53:	{ 1, 2, 6 }	.11...1.	53:	{ 2, 3, 5 }	..11.1..
54:	{ 1, 2, 7 }	.11....1	54:	{ 2, 3, 7 }	..11...1
55:	{ 1, 2, 5 }	.11..1..	55:	{ 2, 3, 6 }	..11..1.
56:	{ 1, 2, 3 }	.111....	56:	{ 2, 3, 4 }	..111...

Figure 6.5-B: Combinations $\binom{8}{3}$ via enup moves (left), and via endo moves (Chase's sequence, right).

the move is homogeneous, no other element is crossed). This property is a direct consequence of the fact that all moves with enup and endo order are by at most two positions. An implementation of an iterative algorithm for the computation of the combinations in enup order is [FXT: `class combination_enup` in comb/combination-enup.h].

```
class combination_enup
{
public:
    ulong *x_;  // combination: k elements 0<=x[j]<k in increasing order
    ulong *s_;  // aux: start of range for enup moves
    ulong *a_;  // aux: actual start position of enup moves
    ulong n_, k_; // Combination (n choose k)
```

```
public:
    combination_enup(ulong n, ulong k)
    {
        n_ = n;
        k_ = k;
        x_ = new ulong[k_+1];  // incl. high sentinel
        s_ = new ulong[k_+1];  // incl. high sentinel
        a_ = new ulong[k_];
        x_[k_] = n_;
        first();
    }

  [--snip--]

    void first()
    {
        for (ulong j=0; j<k_; ++j)  x_[j] = j;
        for (ulong j=0; j<k_; ++j)  s_[j] = j;
        for (ulong j=0; j<k_; ++j)  a_[j] = x_[j];
    }

```

The successor of the current combination is computed by finding the range of possible movements (variable `m`) and, unless the range is empty, move until we are back at the start position:

```
    ulong next()
    // Return position where track changed, return k with last combination
    {
        ulong j = k_;
        while ( j-- )  // loop over tracks
        {
            const ulong sj = s_[j];
            const ulong m = x_[j+1] - sj - 1;

            if ( 0!=m )  // unless range empty
            {
                ulong u = x_[j] - sj;

                // move right on even positions:
                if ( 0==(sj&1) )  u = next_enup(u, m);
                else              u = next_endo(u, m);

                u += sj;

                if ( u != a_[j] )  // next pos != start position
                {
                    x_[j] = u;
                    s_[j+1] = u+1;
                    return j;
                }
            }

            a_[j] = x_[j];
        }

        return  k_;  // current combination is last
    }
};
```

The routine is similar to the one used for the Eades-McKay sequence that is given on page 184. The combinations $\binom{32}{20}$ are generated at a rate of 45 million objects per second, the combinations $\binom{32}{12}$ at a

rate of 55 million per second.

The only change in the implementation for computing the endo ordering is (at the obvious place in the code) [FXT: comb/combination-endo.h]:

```
                // move right on odd positions:
                if ( 0==(sj&1) )  u = next_endo(u, m);
                else              u = next_enup(u, m);
```

The ordering obtained by endo moves is called *Chase's sequence*. Figure 6.5-B was created with the programs [FXT: comb/combination-enup-demo.cc] and [FXT: comb/combination-endo-demo.cc].

The underlying recursion for the list $U(n,k)$ of combinations $\binom{n}{k}$ in enup order is

$$U(n,k) \;=\; \begin{matrix}[(n)\,.\,(n-1)\,.\,U(n-2,\,k-2)] \\ [(n)\,.\,U(n-2,\,k-1) \qquad\quad] \\ [U^{\mathbf{R}}(n-1,\,k) \qquad\quad]\end{matrix} \;=\; \begin{matrix}[1\,1\,.\,U(n-2,\,k-2)] \\ [1\,0\,.\,U(n-2,\,k-1)] \\ [0\,.\,U^{\mathbf{R}}(n-1,\,k) \quad]\end{matrix} \tag{6.5-1}$$

The recursion is very similar to relation 6.4-1 on page 183, therefore a recursive routine is easy to obtain. The crucial part of the routine is [FXT: comb/combination-enup-rec-demo.cc]:

```
void
comb_enup(ulong n, ulong k, bool z)
{
    if ( k==n )  { visit();  return; }

    if ( z )  // forward:
    {
        if ( (n>=2) && (k>=2) )  { rv[k] = n;  rv[k-1] = n-1;  comb_enup(n-2, k-2, z); }
        if ( (n>=2) && (k>=1) )  { rv[k] = n;  comb_enup(n-2, k-1, z); }
        if ( (n>=1) )            { comb_enup(n-1, k, !z); }
    }
    else     // backward:
    {
        if ( (n>=1) )            { comb_enup(n-1, k, !z); }
        if ( (n>=2) && (k>=1) )  { rv[k] = n;  comb_enup(n-2, k-1, z); }
        if ( (n>=2) && (k>=2) )  { rv[k] = n;  rv[k-1] = n-1;  comb_enup(n-2, k-2, z); }
    }
}
```

A recursion for the complement sequence (with respect to the delta sets) is

$$U'(n,k) \;=\; \begin{matrix}[(n)\,.\,U'^{\mathbf{R}}(n-1,\,k-1) \quad] \\ [(n-1)\,.\,U'(n-2,\,k-1)] \\ [U'(n-2,\,k) \qquad\quad]\end{matrix} \;=\; \begin{matrix}[1\,.\,U'^{\mathbf{R}}(n-1,\,k-1)] \\ [0\,1\,.\,U'(n-2,\,k-1)] \\ [0\,0\,.\,U'(n-2,\,k) \quad]\end{matrix} \tag{6.5-2}$$

The condition for the recursion end has to be modified:

```
void
comb_enup_compl(ulong n, ulong k, bool z)
{
    if ( (k==0) || (k==n) )  { visit();  return; }

    if ( z )  // forward:
    {
        if ( (n>=1) && (k>=1) )  { rv[k] = n;  comb_enup_compl(n-1, k-1, !z); }   // 1
        if ( (n>=2) && (k>=1) )  { rv[k] = n-1;  comb_enup_compl(n-2, k-1, z); }  // 01
        if ( (n>=2) )            { comb_enup_compl(n-2, k-0, z); }                // 00
    }
    else     // backward:
    {
        if ( (n>=2) )            { comb_enup_compl(n-2, k-0, z); }                // 00
        if ( (n>=2) && (k>=1) )  { rv[k] = n-1;  comb_enup_compl(n-2, k-1, z); }  // 01
        if ( (n>=1) && (k>=1) )  { rv[k] = n;  comb_enup_compl(n-1, k-1, !z); }   // 1
    }
}
```

The algorithm for Chase's sequence that generates delta sets is described in [170]. An implementation is given in [FXT: `class combination_chase` in comb/combination-chase.h]. The routine generates about 64 million combinations per second both for $\binom{32}{20}$ and $\binom{32}{12}$. See [FXT: comb/combination-chase-demo.cc] for the usage of the class.

6.6 Recursive generation of certain orderings

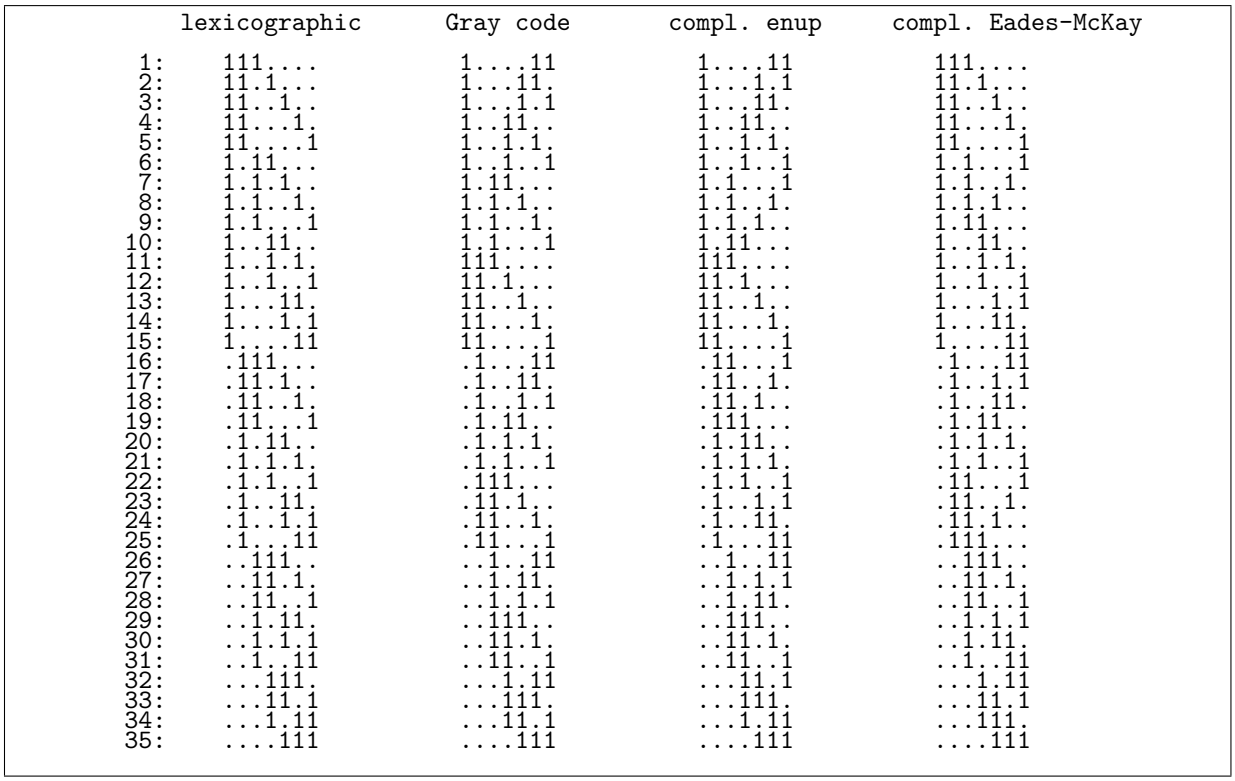

	lexicographic	Gray code	compl. enup	compl. Eades-McKay
1:	111....	1....11	1....11	111....
2:	11.1...	1...11.	1...1.1	11.1...
3:	11..1..	1...1.1	1...11.	11..1..
4:	11...1.	1..11..	1..11..	11...1.
5:	11....1	1..1.1.	1..1.1.	11....1
6:	1.11...	1..1..1	1..1..1	1.1...1
7:	1.1.1..	1.11...	1.1...1	1.1..1.
8:	1.1..1.	1.1.1..	1.1..1.	1.1.1..
9:	1.1...1	1.1..1.	1.1.1..	1.11...
10:	1..11..	1.1...1	1.11...	1..11..
11:	1..1.1.	111....	111....	1..1.1.
12:	1..1..1	11.1...	11.1...	1..1..1
13:	1...11.	11..1..	11..1..	1...1.1
14:	1...1.1	11...1.	11...1.	1...11.
15:	1....11	11....1	11....1	1....11
16:	.111...	.1...11	.11...1	.1...11
17:	.11.1..	.1..11.	.11..1.	.1..1.1
18:	.11..1.	.1..1.1	.11.1..	.1..11.
19:	.11...1	.1.11..	.111...	.1.11..
20:	.1.11..	.1.1.1.	.1.11..	.1.1.1.
21:	.1.1.1.	.1.1..1	.1.1.1.	.1.1..1
22:	.1.1..1	.111...	.1.1..1	.11...1
23:	.1..11.	.11.1..	.1..1.1	.11..1.
24:	.1..1.1	.11..1.	.1..11.	.11.1..
25:	.1...11	.11...1	.1...11	.111...
26:	..111..	..1..11	..1..11	..111..
27:	..11.1.	..1.11.	..1.1.1	..11.1.
28:	..11..1	..1.1.1	..1.11.	..11..1
29:	..1.11.	..111..	..111..	..1.1.1
30:	..1.1.1	..11.1.	..11.1.	..1.11.
31:	..1..11	..11..1	..11..1	..1..11
32:	...111.	...1.11	...11.1	...1.11
33:	...11.1	...111.	...111.	...11.1
34:	...1.11	...11.1	...1.11	...111.
35:	111	111	111	111

Figure 6.6-A: All combinations $\binom{7}{3}$ in lexicographic, minimal-change, complemented enup, and complemented Eades-McKay order (from left to right).

We give a simple recursive routine to obtain the orders shown in figure 6.6-A. The combinations are generated as sets [FXT: `class comb_rec` in comb/combination-rec.h]:

```
class comb_rec
{
public:
    ulong n_, k_; // (n choose k)
    ulong *rv_;   //  combination: k elements 0<=x[j]<k in increasing order
    // == Record of Visits in graph
    ulong rq_;   // condition that determines the order:
    // 0 ==> lexicographic order
    // 1 ==> Gray code
    // 2 ==> complemented enup order
    // 3 ==> complemented Eades-McKay sequence
    ulong nq_;   // whether to reverse order
  [--snip--]
    void (*visit_)(const comb_rec &);  // function to call with each combination
  [--snip--]

    void generate(void (*visit)(const comb_rec &), ulong rq, ulong nq=0)
    {
        visit_ = visit;
        rq_ = rq;
        nq_ = nq;
        ct_ = 0;
        rct_ = 0;
        next_rec(0);
    }
```

The recursion function is given in [FXT: comb/combination-rec.cc]:

```
void comb_rec::next_rec(ulong d)
{
```

```
3       ulong r = k_ - d;  // number of elements remaining
4       if ( 0==r )  visit_(*this);
5       else
6       {
7           ulong rv1 = rv_[d-1];  // left neighbor
8           bool q;
9           switch ( rq_ )
10          {
11          case 0:  q = 1;  break;          //  0 ==> lexicographic order
12          case 1:  q = !(d&1);  break;     //  1 ==> Gray code
13          case 2:  q = rv1&1;  break;      //  2 ==> complemented enup order
14          case 3:  q = (d^rv1)&1;  break;  //  3 ==> complemented Eades-McKay sequence
15          default: q = 1;
16          }
17          q ^= nq_;  // reversed order if  nq == true
18
19          if ( q )  // forward:
20              for (ulong x=rv1+1; x<=n_-r; ++x)  { rv_[d] = x;  next_rec(d+1); }
21          else     // backward:
22              for (ulong x=n_-r; (long)x>=(long)rv1+1; --x)  { rv_[d] = x;  next_rec(d+1); }
23      }
24  }
```

Figure 6.6-A was created with the program [FXT: comb/combination-rec-demo.cc]. The routine generates the combinations $\binom{32}{20}$ at a rate of about 32 million objects per second. The combinations $\binom{32}{12}$ are generated at a rate of 45 million objects per second.

Chapter 7

Compositions

The *compositions* of n into (at most) k parts ('k-compositions of n') are the ordered tuples $(x_0, x_1, \ldots, x_{k-1})$ where $x_0 + x_1 + \ldots + x_{k-1} = n$ and $0 \le x_i \le n$. Order matters: one 4-composition of 7 is $(0, 1, 5, 1)$, different ones are $(5, 0, 1, 1)$ and $(0, 5, 1, 1)$. To obtain the compositions of n into exactly k parts (where $k \le n$) generate the compositions of $n - k$ into k parts and add one to each position.

7.1 Co-lexicographic order

The compositions in co-lexicographic (colex) order are shown in figure 7.1-A. The implementation is [FXT: `class composition_colex` in comb/composition-colex.h]:

```
1   class composition_colex
2   {
3   public:
4       ulong n_, k_;  // composition of n into k parts
5       ulong *x_;     // data (k elements)
6     [--snip--]
7
8       void first()
9       {
10          x_[0] = n_;  // all in first position
11          for (ulong k=1; k<k_; ++k)  x_[k] = 0;
12      }
13
14      void last()
15      {
16          for (ulong k=0; k<k_; ++k)  x_[k] = 0;
17          x_[k_-1] = n_;  // all in last position
18      }
19    [--snip--]
```

The methods to compute the successor and predecessor are:

```
1       ulong next()
2       // Return position of rightmost change, return k with last composition.
3       {
4           ulong j = 0;
5           while ( 0==x_[j] )  ++j;  // find first nonzero
6
7           if ( j==k_-1 )  return k_;  // current composition is last
8
9           ulong v = x_[j];  // value of first nonzero
10          x_[j] = 0;        // set to zero
11          x_[0] = v - 1;    // value-1 to first position
12          ++j;
13          ++x_[j];          // increment next position
14
15          return  j;
```

	composition	chg	combination
1:	[3 . . . .]	4	111....
2:	[2 1 . . .]	1	11.1...
3:	[1 2 . . .]	1	1.11...
4:	[. 3 . . .]	1	.111...
5:	[2 . 1 . .]	2	11..1..
6:	[1 1 1 . .]	1	1.1.1..
7:	[. 2 1 . .]	1	.11.1..
8:	[1 . 2 . .]	2	1..11..
9:	[. 1 2 . .]	1	.1.11..
10:	[. . 3 . .]	2	..111..
11:	[2 . . 1 .]	3	11...1.
12:	[1 1 . 1 .]	1	1.1..1.
13:	[. 2 . 1 .]	1	.11..1.
14:	[1 . 1 1 .]	2	1..1.1.
15:	[. 1 1 1 .]	1	.1.1.1.
16:	[. . 2 1 .]	2	..11.1.
17:	[1 . . 2 .]	3	1...11.
18:	[. 1 . 2 .]	1	.1..11.
19:	[. . 1 2 .]	2	..1.11.
20:	[. . . 3 .]	3	...111.
21:	[2 . . . 1]	4	11....1
22:	[1 1 . . 1]	1	1.1...1
23:	[. 2 . . 1]	1	.11...1
24:	[1 . 1 . 1]	2	1..1..1
25:	[. 1 1 . 1]	1	.1.1..1
26:	[. . 2 . 1]	2	..11..1
27:	[1 . . 1 1]	3	1...1.1
28:	[. 1 . 1 1]	1	.1..1.1
29:	[. . 1 1 1]	2	..1.1.1
30:	[. . . 2 1]	3	...11.1
31:	[1 . . . 2]	4	1....11
32:	[. 1 . . 2]	1	.1...11
33:	[. . 1 . 2]	2	..1..11
34:	[. . . 1 2]	3	...1.11
35:	[. . . . 3]	4	111

	composition	chg	combination
1:	[7 . .]	2	1111111..
2:	[6 1 .]	1	111111.1.
3:	[5 2 .]	1	11111.11.
4:	[4 3 .]	1	1111.111.
5:	[3 4 .]	1	111.1111.
6:	[2 5 .]	1	11.11111.
7:	[1 6 .]	1	1.111111.
8:	[. 7 .]	1	.1111111.
9:	[6 . 1]	2	111111..1
10:	[5 1 1]	1	11111.1.1
11:	[4 2 1]	1	1111.11.1
12:	[3 3 1]	1	111.111.1
13:	[2 4 1]	1	11.1111.1
14:	[1 5 1]	1	1.11111.1
15:	[. 6 1]	1	.111111.1
16:	[5 . 2]	2	11111..11
17:	[4 1 2]	1	1111.1.11
18:	[3 2 2]	1	111.11.11
19:	[2 3 2]	1	11.111.11
20:	[1 4 2]	1	1.1111.11
21:	[. 5 2]	1	.11111.11
22:	[4 . 3]	2	1111..111
23:	[3 1 3]	1	111.1.111
24:	[2 2 3]	1	11.11.111
25:	[1 3 3]	1	1.111.111
26:	[. 4 3]	1	.1111.111
27:	[3 . 4]	2	111..1111
28:	[2 1 4]	1	11.1.1111
29:	[1 2 4]	1	1.11.1111
30:	[. 3 4]	1	.111.1111
31:	[2 . 5]	2	11..11111
32:	[1 1 5]	1	1.1.11111
33:	[. 2 5]	1	.11.11111
34:	[1 . 6]	2	1..111111
35:	[. 1 6]	1	.1.111111
36:	[. . 7]	2	..1111111

Figure 7.1-A: The compositions of 3 into 5 parts in co-lexicographic order and positions of the rightmost change, and delta sets of the corresponding combinations (left); and the corresponding data for compositions of 7 into 3 parts (right). Dots denote zeros.

```
16      }
17
18      ulong prev()
19      // Return position of rightmost change, return k with last composition.
20      {
21          const ulong v = x_[0];   // value at first position
22
23          if ( n_==v )  return k_; // current composition is first
24
25          x_[0] = 0;          // set first position to zero
26          ulong j = 1;
27          while ( 0==x_[j] )  ++j;  // find next nonzero
28          --x_[j];                  // decrement value
29          x_[j-1] = 1 + v;          // set previous position
30
31          return  j;
32      }
```

With each transition at most 3 entries are changed. The compositions of 10 into 30 parts (sparse case) are generated at a rate of about 110 million per second, the compositions of 30 into 10 parts (dense case) at about 200 million per second [FXT: comb/composition-colex-demo.cc]. With the dense case (corresponding to the right of figure 7.1-A) the computation is faster as the position to change is found earlier.

Optimized implementation

An implementation that is efficient also for the sparse case (that is, k much greater than n) is [FXT: `class composition_colex2` in comb/composition-colex2.h]. One additional variable `p0` records the position of the first nonzero entry. The method to compute the successor is:

```
1   class composition_colex2
2   {
3     [--snip--]
4       ulong next()
5       // Return position of rightmost change, return k with last composition.
6       {
7           ulong j = p0_;  // position of first nonzero
8
9           if ( j==k_-1 )  return k_;  // current composition is last
10
11          ulong v = x_[j];  // value of first nonzero
12          x_[j] = 0;        // set to zero
13          --v;
14          x_[0] = v;        // value-1 to first position
15
16          ++p0_;            // first nonzero one more right except ...
17          if ( 0!=v )  p0_ = 0;   // ... if value v was not one
18
19          ++j;
20          ++x_[j];          // increment next position
21
22          return  j;
23      }
24  };
```

About 182 million compositions are generated per second, independent of either n and k [FXT: comb/composition-colex2-demo.cc]. With the line

```
#define COMP_COLEX2_MAX_ARRAY_LEN  128
```

just before the class definition an array is used instead of a pointer. The fixed array length limits the value of k so by default the line is commented out. Using an array gives great speedup, the rate is about 365 million per second (about 6 CPU cycles per update).

7.2 Co-lexicographic order for compositions into exactly k parts

The compositions of n into exactly k parts (where $k \geq n$) can be obtained from the compositions of $n - k$ into at most k parts as shown in figure 7.2-A. The listing was created with the program [FXT: comb/composition-ex-colex-demo.cc]. The compositions can be generated in co-lexicographic order using [FXT: `class composition_ex_colex` in comb/composition-ex-colex.h]:

```
1   class composition_ex_colex
2   {
3   public:
4       ulong n_, k_;  // composition of n into exactly k parts
5       ulong *x_;     // data (k elements)
6       ulong nk1_;    // ==n-k+1
7
8   public:
9       composition_ex_colex(ulong n, ulong k)
10      // Must have n>=k
11      {
12          n_ = n;
13          k_ = k;
14          nk1_ = n - k + 1;  // must be >= 1
15          if ( (long)nk1_ < 1 )  nk1_ = 1;  // avoid hang with invalid pair n,k
16          x_ = new ulong[k_ + 1];
17          x_[k] = 0;  // not one
18          first();
19      }
20    [--snip--]
```

The variable `nk1_` is the maximal entry in the compositions:

```
              exact comp.  chg      composition
 1:     [ 4 1 1 1 1 ]  4      [ 3 . . . . ]
 2:     [ 3 2 1 1 1 ]  1      [ 2 1 . . . ]
 3:     [ 2 3 1 1 1 ]  1      [ 1 2 . . . ]
 4:     [ 1 4 1 1 1 ]  1      [ . 3 . . . ]
 5:     [ 3 1 2 1 1 ]  2      [ 2 . 1 . . ]
 6:     [ 2 2 2 1 1 ]  1      [ 1 1 1 . . ]
 7:     [ 1 3 2 1 1 ]  1      [ . 2 1 . . ]
 8:     [ 2 1 3 1 1 ]  2      [ 1 . 2 . . ]
 9:     [ 1 2 3 1 1 ]  1      [ . 1 2 . . ]
10:     [ 1 1 4 1 1 ]  2      [ . . 3 . . ]
11:     [ 3 1 1 2 1 ]  3      [ 2 . . 1 . ]
12:     [ 2 2 1 2 1 ]  1      [ 1 1 . 1 . ]
13:     [ 1 3 1 2 1 ]  1      [ . 2 . 1 . ]
14:     [ 2 1 2 2 1 ]  2      [ 1 . 1 1 . ]
15:     [ 1 2 2 2 1 ]  1      [ . 1 1 1 . ]
16:     [ 1 1 3 2 1 ]  2      [ . . 2 1 . ]
17:     [ 2 1 1 3 1 ]  3      [ 1 . . 2 . ]
18:     [ 1 2 1 3 1 ]  1      [ . 1 . 2 . ]
19:     [ 1 1 2 3 1 ]  2      [ . . 1 2 . ]
20:     [ 1 1 1 4 1 ]  3      [ . . . 3 . ]
21:     [ 3 1 1 1 2 ]  4      [ 2 . . . 1 ]
22:     [ 2 2 1 1 2 ]  1      [ 1 1 . . 1 ]
23:     [ 1 3 1 1 2 ]  1      [ . 2 . . 1 ]
24:     [ 2 1 2 1 2 ]  2      [ 1 . 1 . 1 ]
25:     [ 1 2 2 1 2 ]  1      [ . 1 1 . 1 ]
26:     [ 1 1 3 1 2 ]  2      [ . . 2 . 1 ]
27:     [ 2 1 1 2 2 ]  3      [ 1 . . 1 1 ]
28:     [ 1 2 1 2 2 ]  1      [ . 1 . 1 1 ]
29:     [ 1 1 2 2 2 ]  2      [ . . 1 1 1 ]
30:     [ 1 1 1 3 2 ]  3      [ . . . 2 1 ]
31:     [ 2 1 1 1 3 ]  4      [ 1 . . . 2 ]
32:     [ 1 2 1 1 3 ]  1      [ . 1 . . 2 ]
33:     [ 1 1 2 1 3 ]  2      [ . . 1 . 2 ]
34:     [ 1 1 1 2 3 ]  3      [ . . . 1 2 ]
35:     [ 1 1 1 1 4 ]  4      [ . . . . 3 ]
```

Figure 7.2-A: The compositions of $n = 8$ into exactly $k = 5$ parts (left) are obtained from the compositions of $n - k = 3$ into (at most) $k = 5$ parts (right). Co-lexicographic order, dots denote zeros.

```
1       void first()
2       {
3           x_[0] = nk1_;  // all in first position
4           for (ulong k=1; k<k_; ++k)  x_[k] = 1;
5       }
6
7       void last()
8       {
9           for (ulong k=0; k<k_; ++k)  x_[k] = 1;
10          x_[k_-1] = nk1_;  // all in last position
11      }
```

The methods for computing the successor and predecessor are adaptations from the routines from the compositions into at most k parts:

```
1       ulong next()
2       // Return position of rightmost change, return k with last composition.
3       {
4           ulong j = 0;
5           while ( 1==x_[j] )  ++j;  // find first greater than one
6
7           if ( j==k_ )  return k_;  // current composition is last
8
9           ulong v = x_[j];  // value of first greater one
10          x_[j] = 1;        // set to one
11          x_[0] = v - 1;    // value-1 to first position
12          ++j;
13          ++x_[j];          // increment next position
14
15          return  j;
16      }
17
```

```
18      ulong prev()
19      // Return position of rightmost change, return k with last composition.
20      {
21          const ulong v = x_[0];   // value at first position
22
23          if ( nk1_==v )  return k_; // current composition is first
24
25          x_[0] = 1;          // set first position to one
26          ulong j = 1;
27          while ( 1==x_[j] )  ++j;  // find next greater than one
28          --x_[j];                  // decrement value
29          x_[j-1] = 1 + v;          // set previous position
30
31          return  j;
32      }
33  };
```

The routines are as fast as the generation into at most k parts with the corresponding parameters: the compositions of 40 into 10 parts are generated at about 200 million per second.

7.3 Compositions and combinations

```
              combination    delta set       composition
        1:    [ 0 1 2 ]        111...        [ 3 . . . ]
        2:    [ 0 2 3 ]        1.11..        [ 1 2 . . ]
        3:    [ 1 2 3 ]        .111..        [ . 3 . . ]
        4:    [ 0 1 3 ]        11.1..        [ 2 1 . . ]
        5:    [ 0 3 4 ]        1..11.        [ 1 . 2 . ]
        6:    [ 1 3 4 ]        .1.11.        [ . 1 2 . ]
        7:    [ 2 3 4 ]        ..111.        [ . . 3 . ]
        8:    [ 0 2 4 ]        1.1.1.        [ 1 1 1 . ]
        9:    [ 1 2 4 ]        .11.1.        [ . 2 1 . ]
       10:    [ 0 1 4 ]        11..1.        [ 2 . 1 . ]
       11:    [ 0 4 5 ]        1...11        [ 1 . . 2 ]
       12:    [ 1 4 5 ]        .1..11        [ . 1 . 2 ]
       13:    [ 2 4 5 ]        ..1.11        [ . . 1 2 ]
       14:    [ 3 4 5 ]        ...111        [ . . . 3 ]
       15:    [ 0 3 5 ]        1..1.1        [ 1 . 1 1 ]
       16:    [ 1 3 5 ]        .1.1.1        [ . 1 1 1 ]
       17:    [ 2 3 5 ]        ..11.1        [ . . 2 1 ]
       18:    [ 0 2 5 ]        1.1..1        [ 1 1 . 1 ]
       19:    [ 1 2 5 ]        .11..1        [ . 2 . 1 ]
       20:    [ 0 1 5 ]        11...1        [ 2 . . 1 ]
```

Figure 7.3-A: Combinations 6 choose 3 (left) and the corresponding compositions of 3 into 4 parts (right). Note that while the sequence of combinations has a minimal-change property the corresponding sequence of compositions does not. Dots denote zeros.

Figure 7.3-A shows the correspondence between compositions and combinations. The listing was generated using the program [FXT: comb/comb2comp-demo.cc]. Entries in the left column are combinations of 3 parts out of 6. The middle column is the representation of the combinations as delta sets. It also is a binary representation of a composition: A run of r consecutive ones corresponds to an entry r in the composition at the right.

Now write $P(n,k)$ for the compositions of n into (at most) k parts and $C(N,K)$ for the combination $\binom{N}{K}$: A composition of n into at most k parts corresponds to a combination of $K = n$ parts from $N = n+k-1$ elements, symbolically,

$$P(n,\,k) \quad \leftrightarrow \quad C(N,\,K) \;=\; C(n+k-1,\,n) \tag{7.3-1a}$$

A combination of K elements out of N corresponds to a composition of n into at most k parts where $n = K$ and $k = N - K + 1$:

$$C(N,\,K) \quad \leftrightarrow \quad P(n,\,k) \;=\; P(K,\,N-K+1) \tag{7.3-1b}$$

Routines for the conversion between combinations and compositions are given in [FXT: comb/comp2comb.h]. The following routine converts a composition into the corresponding combination:

```
1   inline void comp2comb(const ulong *p, ulong k, ulong *c)
2   // Convert composition P(*, k) in p[] to combination in c[]
3   {
4       for (ulong j=0,i=0,z=0; j<k; ++j)
5       {
6           ulong pj = p[j];
7           for (ulong w=0; w<pj; ++w)   c[i++] = z++;
8           ++z;
9       }
10  }
```

Conversion of a combination into the corresponding composition:

```
1   inline void comb2comp(const ulong *c, ulong N, ulong K, ulong *p)
2   // Convert combination C(N, K) in c[] to composition P(*,k) in p[]
3   // Must have: K>0
4   {
5       ulong k = N-K+1;
6       for (ulong z=0; z<k; ++z)  p[z] = 0;
7       --k;
8       ulong c1 = N;
9       while ( K-- )
10      {
11          ulong c0 = c[K];
12          ulong d = c1 - c0;
13          k -= (d-1);
14          ++p[k];
15          c1 = c0;
16      }
17  }
```

7.4 Minimal-change orders

A minimal-change order (Gray code) for compositions is such that with each transition one entry is increased by one and another is decreased by one. A recursion for the compositions $P(n,k)$ of n into k parts in lexicographic order is

$$P(n,k) \;=\; \begin{matrix} [0\,.\,P(n-0,k-1)] \\ [1\,.\,P(n-1,k-1)] \\ [2\,.\,P(n-2,k-1)] \\ [3\,.\,P(n-3,k-1)] \\ [4\,.\,P(n-4,k-1)] \\ [\quad\vdots\qquad\quad] \\ [n\,.\,P(0,k-1)\quad\;] \end{matrix} \tag{7.4-1}$$

A simple variation gives a Gray code, we change the direction if the element is even:

$$P(n,k) \;=\; \begin{matrix} [0\,.\,P^{\mathbf{R}}(n-0,k-1)] \\ [1\,.\,P(n-1,k-1)\;\;] \\ [2\,.\,P^{\mathbf{R}}(n-2,k-1)] \\ [3\,.\,P(n-3,k-1)\;\;] \\ [4\,.\,P^{\mathbf{R}}(n-4,k-1)] \\ [\quad\vdots\qquad\quad] \end{matrix} \tag{7.4-2}$$

	composition	combination		composition	combination
1:	[. . . 3 .]	...111.	[3 4 5]	[3 . . . .]	111....	[0 1 2]
2:	[. 1 . 2 .]	.1..11.	[1 4 5]	[2 1 . . .]	11.1...	[0 1 3]
3:	[1 . . 2 .]	1...11.	[0 4 5]	[1 2 . . .]	1.11...	[0 2 3]
4:	[. . 1 2 .]	..1.11.	[2 4 5]	[. 3 . . .]	.111...	[1 2 3]
5:	[. . 2 1 .]	..11.1.	[2 3 5]	[. 2 1 . .]	.11.1..	[1 2 4]
6:	[. 1 1 1 .]	.1.1.1.	[1 3 5]	[1 1 1 . .]	1.1.1..	[0 2 4]
7:	[1 . 1 1 .]	1..1.1.	[0 3 5]	[2 . 1 . .]	11..1..	[0 1 4]
8:	[2 . . 1 .]	11...1.	[0 1 5]	[1 . 2 . .]	1..11..	[0 3 4]
9:	[1 1 . 1 .]	1.1..1.	[0 2 5]	[. 1 2 . .]	.1.11..	[1 3 4]
10:	[. 2 . 1 .]	.11..1.	[1 2 5]	[. . 3 . .]	..111..	[2 3 4]
11:	[. 3 . . .]	.111...	[1 2 3]	[. . 2 1 .]	..11.1.	[2 3 5]
12:	[1 2 . . .]	1.11...	[0 2 3]	[1 . 1 1 .]	1..1.1.	[0 3 5]
13:	[2 1 . . .]	11.1...	[0 1 3]	[. 1 1 1 .]	.1.1.1.	[1 3 5]
14:	[3 . . . .]	111....	[0 1 2]	[. 2 . 1 .]	.11..1.	[1 2 5]
15:	[2 . 1 . .]	11..1..	[0 1 4]	[1 1 . 1 .]	1.1..1.	[0 2 5]
16:	[1 1 1 . .]	1.1.1..	[0 2 4]	[2 . . 1 .]	11...1.	[0 1 5]
17:	[. 2 1 . .]	.11.1..	[1 2 4]	[1 . . 2 .]	1...11.	[0 4 5]
18:	[. 1 2 . .]	.1.11..	[1 3 4]	[. 1 . 2 .]	.1..11.	[1 4 5]
19:	[1 . 2 . .]	1..11..	[0 3 4]	[. . 1 2 .]	..1.11.	[2 4 5]
20:	[. . 3 . .]	..111..	[2 3 4]	[. . . 3 .]	...111.	[3 4 5]
21:	[. . 2 . 1]	..11..1	[2 3 6]	[. . . 2 1]	...11.1	[3 4 6]
22:	[. 1 1 . 1]	.1.1..1	[1 3 6]	[1 . . 1 1]	1...1.1	[0 4 6]
23:	[1 . 1 . 1]	1..1..1	[0 3 6]	[. 1 . 1 1]	.1..1.1	[1 4 6]
24:	[2 . . . 1]	11....1	[0 1 6]	[. . 1 1 1]	..1.1.1	[2 4 6]
25:	[1 1 . . 1]	1.1...1	[0 2 6]	[. . 2 . 1]	..11..1	[2 3 6]
26:	[. 2 . . 1]	.11...1	[1 2 6]	[1 . 1 . 1]	1..1..1	[0 3 6]
27:	[. 1 . 1 1]	.1..1.1	[1 4 6]	[. 1 1 . 1]	.1.1..1	[1 3 6]
28:	[1 . . 1 1]	1...1.1	[0 4 6]	[. 2 . . 1]	.11...1	[1 2 6]
29:	[. . 1 1 1]	..1.1.1	[2 4 6]	[1 1 . . 1]	1.1...1	[0 2 6]
30:	[. . . 2 1]	...11.1	[3 4 6]	[2 . . . 1]	11....1	[0 1 6]
31:	[. . . 1 2]	...1.11	[3 5 6]	[1 . . . 2]	1....11	[0 5 6]
32:	[. 1 . . 2]	.1...11	[1 5 6]	[. 1 . . 2]	.1...11	[1 5 6]
33:	[1 . . . 2]	1....11	[0 5 6]	[. . 1 . 2]	..1..11	[2 5 6]
34:	[. . 1 . 2]	..1..11	[2 5 6]	[. . . 1 2]	...1.11	[3 5 6]
35:	[. . . . 3]	111	[4 5 6]	[. . . . 3]	111	[4 5 6]

Figure 7.4-A: Compositions of 3 into 5 parts and the corresponding combinations as delta sets and sets in two minimal-change orders: order with enup moves (left) and order with modulo moves (right). The ordering by enup moves is a two-close Gray code. Dots denote zeros.

The ordering is shown in figure 7.4-A (left), the corresponding combinations are in the (reversed) enup order from section 6.5.2 on page 186. If we change directions at the odd elements

$$P(n,k) \quad = \quad \begin{matrix} [0 \,.\, P(n-0,k-1) \;\;] \\ [1 \,.\, P^{\mathbf{R}}(n-1,k-1)] \\ [2 \,.\, P(n-2,k-1) \;\;] \\ [3 \,.\, P^{\mathbf{R}}(n-3,k-1)] \\ [4 \,.\, P(n-4,k-1) \;\;] \\ [\quad \vdots \qquad\qquad\;\;] \end{matrix} \tag{7.4-3}$$

then we obtain an ordering (right of figure 7.4-A) corresponding to the combinations are in the (reversed) Eades-McKay order from section 6.4 on page 182. The listings were created with the program [FXT: comb/composition-gray-rec-demo.cc].

Gray codes for combinations correspond to Gray codes for combinations where no element in the delta set crosses another. The standard Gray code for combinations does not lead to a Gray code for compositions as shown in figure 7.3-A on page 197. When the directions in the recursions are always changed then the compositions correspond to combinations that have the complemented delta sets of the standard Gray code in reversed order.

Orderings where the changes involve just one pair of adjacent entries (shown in figure 7.4-B) correspond to the complemented strong Gray codes for combinations. The amount of change is greater than one in general. The listings were created with the program [FXT: comb/combination-rec-demo.cc], see section 6.6

	combination		composition
1:	[0 5 6]	1....11	[1 . . . 2]
2:	[0 4 6]	1...1.1	[1 . . 1 1]
3:	[0 4 5]	1...11.	[1 . . 2 .]
4:	[0 3 4]	1..11..	[1 . 2 . .]
5:	[0 3 5]	1..1.1.	[1 . 1 1 .]
6:	[0 3 6]	1..1..1	[1 . 1 . 1]
7:	[0 2 6]	1.1...1	[1 1 . . 1]
8:	[0 2 5]	1.1..1.	[1 1 . 1 .]
9:	[0 2 4]	1.1.1..	[1 1 1 . .]
10:	[0 2 3]	1.11...	[1 2 . . .]
11:	[0 1 2]	111....	[3 . . . .]
12:	[0 1 3]	11.1...	[2 1 . . .]
13:	[0 1 4]	11..1..	[2 . 1 . .]
14:	[0 1 5]	11...1.	[2 . . 1 .]
15:	[0 1 6]	11....1	[2 . . . 1]
16:	[1 2 6]	.11...1	[. 2 . . 1]
17:	[1 2 5]	.11..1.	[. 2 . 1 .]
18:	[1 2 4]	.11.1..	[. 2 1 . .]
19:	[1 2 3]	.111...	[. 3 . . .]
20:	[1 3 4]	.1.11..	[. 1 2 . .]
21:	[1 3 5]	.1.1.1.	[. 1 1 1 .]
22:	[1 3 6]	.1.1..1	[. 1 1 . 1]
23:	[1 4 6]	.1..1.1	[. 1 . 1 1]
24:	[1 4 5]	.1..11.	[. 1 . 2 .]
25:	[1 5 6]	.1...11	[. 1 . . 2]
26:	[2 5 6]	..1..11	[. . 1 . 2]
27:	[2 4 6]	..1.1.1	[. . 1 1 1]
28:	[2 4 5]	..1.11.	[. . 1 2 .]
29:	[2 3 4]	..111..	[. . 3 . .]
30:	[2 3 5]	..11.1.	[. . 2 1 .]
31:	[2 3 6]	..11..1	[. . 2 . 1]
32:	[3 4 6]	...11.1	[. . . 2 1]
33:	[3 4 5]	...111.	[. . . 3 .]
34:	[3 5 6]	...1.11	[. . . 1 2]
35:	[4 5 6]	111	[. . . . 3]

	combination		composition
1:	[0 1 2]	111....	[3 . . . .]
2:	[0 1 3]	11.1...	[2 1 . . .]
3:	[0 1 4]	11..1..	[2 . 1 . .]
4:	[0 1 5]	11...1.	[2 . . 1 .]
5:	[0 1 6]	11....1	[2 . . . 1]
6:	[0 2 6]	1.1...1	[1 1 . . 1]
7:	[0 2 5]	1.1..1.	[1 1 . 1 .]
8:	[0 2 4]	1.1.1..	[1 1 1 . .]
9:	[0 2 3]	1.11...	[1 2 . . .]
10:	[0 3 4]	1..11..	[1 . 2 . .]
11:	[0 3 5]	1..1.1.	[1 . 1 1 .]
12:	[0 3 6]	1..1..1	[1 . 1 . 1]
13:	[0 4 6]	1...1.1	[1 . . 1 1]
14:	[0 4 5]	1...11.	[1 . . 2 .]
15:	[0 5 6]	1....11	[1 . . . 2]
16:	[1 5 6]	.1...11	[. 1 . . 2]
17:	[1 4 6]	.1..1.1	[. 1 . 1 1]
18:	[1 4 5]	.1..11.	[. 1 . 2 .]
19:	[1 3 4]	.1.11..	[. 1 2 . .]
20:	[1 3 5]	.1.1.1.	[. 1 1 1 .]
21:	[1 3 6]	.1.1..1	[. 1 1 . 1]
22:	[1 2 6]	.11...1	[. 2 . . 1]
23:	[1 2 5]	.11..1.	[. 2 . 1 .]
24:	[1 2 4]	.11.1..	[. 2 1 . .]
25:	[1 2 3]	.111...	[. 3 . . .]
26:	[2 3 4]	..111..	[. . 3 . .]
27:	[2 3 5]	..11.1.	[. . 2 1 .]
28:	[2 3 6]	..11..1	[. . 2 . 1]
29:	[2 4 6]	..1.1.1	[. . 1 1 1]
30:	[2 4 5]	..1.11.	[. . 1 2 .]
31:	[2 5 6]	..1..11	[. . 1 . 2]
32:	[3 5 6]	...1.11	[. . . 1 2]
33:	[3 4 6]	...11.1	[. . . 2 1]
34:	[3 4 5]	...111.	[. . . 3 .]
35:	[4 5 6]	111	[. . . . 3]

Figure 7.4-B: The (reversed) complemented enup ordering (left) and Eades-McKay sequence (right) for combinations correspond to compositions where only two adjacent entries change with each transition, but by more than one in general.

on page 190.

Chapter 8

Subsets

This section gives algorithms to generate all subsets of a set of n elements. There are 2^n subsets, including the empty set.

8.1 Lexicographic order

```
 1:    1....   {0}                  1....   {0}
 2:    11...   {0, 1}               .1...   {1}
 3:    111..   {0, 1, 2}            11...   {0, 1}
 4:    1111.   {0, 1, 2, 3}         ..1..   {2}
 5:    11111   {0, 1, 2, 3, 4}      1.1..   {0, 2}
 6:    111.1   {0, 1, 2, 4}         .11..   {1, 2}
 7:    11.1.   {0, 1, 3}            111..   {0, 1, 2}
 8:    11.11   {0, 1, 3, 4}         ...1.   {3}
 9:    11..1   {0, 1, 4}            1..1.   {0, 3}
10:    1.1..   {0, 2}               .1.1.   {1, 3}
11:    1.11.   {0, 2, 3}            11.1.   {0, 1, 3}
12:    1.111   {0, 2, 3, 4}         ..11.   {2, 3}
13:    1.1.1   {0, 2, 4}            1.11.   {0, 2, 3}
14:    1..1.   {0, 3}               .111.   {1, 2, 3}
15:    1..11   {0, 3, 4}            1111.   {0, 1, 2, 3}
16:    1...1   {0, 4}               ....1   {4}
17:    .1...   {1}                  1...1   {0, 4}
18:    .11..   {1, 2}               .1..1   {1, 4}
19:    .111.   {1, 2, 3}            11..1   {0, 1, 4}
20:    .1111   {1, 2, 3, 4}         ..1.1   {2, 4}
21:    .11.1   {1, 2, 4}            1.1.1   {0, 2, 4}
22:    .1.1.   {1, 3}               .11.1   {1, 2, 4}
23:    .1.11   {1, 3, 4}            111.1   {0, 1, 2, 4}
24:    .1..1   {1, 4}               ...11   {3, 4}
25:    ..1..   {2}                  1..11   {0, 3, 4}
26:    ..11.   {2, 3}               .1.11   {1, 3, 4}
27:    ..111   {2, 3, 4}            11.11   {0, 1, 3, 4}
28:    ..1.1   {2, 4}               ..111   {2, 3, 4}
29:    ...1.   {3}                  1.111   {0, 2, 3, 4}
30:    ...11   {3, 4}               .1111   {1, 2, 3, 4}
31:    ....1   {4}                  11111   {0, 1, 2, 3, 4}
```

Figure 8.1-A: Non-empty subsets of a five element set in lexicographic order for the sets (left), and in lexicographic order for the delta sets (right).

The (nonempty) subsets of a set of five elements in lexicographic order are shown in figure 8.1-A. Note that the lexicographic order is different for the set and the delta set representation.

8.1.1 Generation as delta sets

The listing on the right side of figure 8.1-A is with respect to the delta sets. It was created with the program [FXT: comb/subset-deltalex-demo.cc] which uses [FXT: `class subset_deltalex` in comb/subset-deltalex.h]. The algorithm for the computation of the successor is binary counting:

```
1   class subset_deltalex
2   {
3   public:
4       ulong *d_;  // subset as delta set
5       ulong n_;   // subsets of the  n-set {0,1,2,...,n-1}
6
7   public:
8       subset_deltalex(ulong n)
9       {
10          n_ = n;
11          d_ = new ulong[n+1];
12          d_[n] = 0;  // sentinel
13          first();
14      }
15
16      ~subset_deltalex()
17      { delete [] d_; }
18
19      void first()
20      { for (ulong k=0; k<n_; ++k)  d_[k] = 0; }
21
22      bool next()
23      {
24          ulong k = 0;
25          while ( d_[k]==1 )  { d_[k]=0;  ++k; }
26
27          if ( k==n_ )  return false;
28          else
29          {
30              d_[k] = 1;
31              return true;
32          }
33      }
34
35      const ulong * data()  const { return d_; }
36  };
```

About 180 million subsets per second are generated. A bit-level algorithm to compute the subsets in lexicographic order is given in section 1.28 on page 71.

8.1.2 Generation as sets

The lexicographic order with respect to the set representation is shown at the left side of figure 8.1-A. The routines in [FXT: `class subset_lex` in comb/subset-lex.h] compute the non-empty sets in the corresponding representation:

```
1   class subset_lex
2   {
3   public:
4       ulong *x_;  // subset of {0,1,2,...,n-1}
5       ulong n_;   // number of elements in set
6       ulong k_;   // index of last element in subset
7       // Number of elements in subset == k+1
8
9   public:
10      subset_lex(ulong n)
11      {
12          n_ = n;
13          x_ = new ulong[n_];
14          first();
15      }
16
17      ~subset_lex()  { delete [] x_; }
18
19      ulong first()
20      {
21          k_ = 0;
22          x_[0] = 0;
```

```
23          return  k_ + 1;
24      }
25
26      ulong last()
27      {
28          k_ = 0;
29          x_[0] = n_ - 1;
30          return  k_ + 1;
31      }
32    [--snip--]
```

The methods `next()` and `prev()` compute the successor and predecessor, respectively:

```
1       ulong next()
2       // Generate next subset
3       // Return number of elements in subset
4       // Return zero if current == last
5       {
6           if ( x_[k_] == n_-1 )  // last element is max ?
7           {
8               if ( k_==0 )  { first();  return 0; }
9
10              --k_;      // remove last element
11              x_[k_]++;  // increase last element
12          }
13          else  // add next element from set:
14          {
15              ++k_;
16              x_[k_] = x_[k_-1] + 1;
17          }
18
19          return  k_ + 1;
20      }
21
22      ulong prev()
23      // Generate previous subset
24      // Return number of elements in subset
25      // Return zero if current == first
26      {
27          if ( k_ == 0 )  // only one element ?
28          {
29              if ( x_[0]==0 )  { last();  return 0; }
30
31              x_[0]--;  // decr first element
32              x_[++k_] = n_ - 1;     // add element
33          }
34          else
35          {
36              if ( x_[k_] == x_[k_-1]+1 )  --k_;  // remove last element
37              else
38              {
39                  x_[k_]--;  // decr last element
40                  x_[++k_] = n_ - 1;     // add element
41              }
42          }
43
44          return  k_ + 1;
45      }
46
47      const ulong * data()  const { return x_; }
48  };
```

More than 235 million subsets per second are generated [FXT: comb/subset-lex-demo.cc]. A generalization with mixed radix numbers is described in section 9.3 on page 225.

8.2 Minimal-change order

8.2.1 Generation as delta sets

The subsets of a set with 5 elements in minimal-change order are shown in figure 8.2-A. The implementation [FXT: `class subset_gray_delta` in comb/subset-gray-delta.h] uses the Gray code of binary words and updates the position corresponding to the bit that changes in the Gray code:

```
1   class subset_gray_delta
```

```
 0:  .....  {}                          0:  11111  { 0, 1, 2, 3, 4 }
 1:  1....  {0}                         1:  .1111  { 1, 2, 3, 4 }
 2:  11...  {0, 1}                      2:  ..111  { 2, 3, 4 }
 3:  .1...  {1}                         3:  1.111  { 0, 2, 3, 4 }
 4:  .11..  {1, 2}                      4:  1..11  { 0, 3, 4 }
 5:  111..  {0, 1, 2}                   5:  ...11  { 3, 4 }
 6:  1.1..  {0, 2}                      6:  .1.11  { 1, 3, 4 }
 7:  ..1..  {2}                         7:  11.11  { 0, 1, 3, 4 }
 8:  ..11.  {2, 3}                      8:  11..1  { 0, 1, 4 }
 9:  1.11.  {0, 2, 3}                   9:  .1..1  { 1, 4 }
10:  1111.  {0, 1, 2, 3}               10:  ....1  { 4 }
11:  .111.  {1, 2, 3}                  11:  1...1  { 0, 4 }
12:  .1.1.  {1, 3}                     12:  1.1.1  { 0, 2, 4 }
13:  11.1.  {0, 1, 3}                  13:  ..1.1  { 2, 4 }
14:  1..1.  {0, 3}                     14:  .11.1  { 1, 2, 4 }
15:  ...1.  {3}                        15:  111.1  { 0, 1, 2, 4 }
16:  ...11  {3, 4}                     16:  111..  { 0, 1, 2 }
17:  1..11  {0, 3, 4}                  17:  .11..  { 1, 2 }
18:  11.11  {0, 1, 3, 4}               18:  ..1..  { 2 }
19:  .1.11  {1, 3, 4}                  19:  1.1..  { 0, 2 }
20:  .1111  {1, 2, 3, 4}               20:  1....  { 0 }
21:  11111  {0, 1, 2, 3, 4}            21:  .....  {  }
22:  1.111  {0, 2, 3, 4}               22:  .1...  { 1 }
23:  ..111  {2, 3, 4}                  23:  11...  { 0, 1 }
24:  ..1.1  {2, 4}                     24:  11.1.  { 0, 1, 3 }
25:  1.1.1  {0, 2, 4}                  25:  .1.1.  { 1, 3 }
26:  111.1  {0, 1, 2, 4}               26:  ...1.  { 3 }
27:  .11.1  {1, 2, 4}                  27:  1..1.  { 0, 3 }
28:  .1..1  {1, 4}                     28:  1.11.  { 0, 2, 3 }
29:  11..1  {0, 1, 4}                  29:  ..11.  { 2, 3 }
30:  1...1  {0, 4}                     30:  .111.  { 1, 2, 3 }
31:  ....1  {4}                        31:  1111.  { 0, 1, 2, 3 }
```

Figure 8.2-A: The subsets of the set {0, 1, 2, 3, 4} in minimal-change order (left), and complemented minimal-change order (right). The changes are on the same places for both orders.

```
2   // Subsets of the set {0,1,2,...,n-1} in minimal-change (Gray code) order.
3   {
4   public:
5       ulong *x_;  // current subset as delta-set
6       ulong n_;   // number of elements in set <= BITS_PER_LONG
7       ulong j_;   // position of last change
8       ulong ct_;  // gray_code(ct_) corresponds to the current subset
9       ulong mct_; // max value of ct.
10
11  public:
12      subset_gray_delta(ulong n)
13      {
14          n_ = (n ? n : 1);  // not zero
15          x_ = new ulong[n_];
16          mct_ = (1UL<<n) - 1;
17          first(0);
18      }
19
20      ~subset_gray_delta()  { delete [] x_; }
21
```

In the initializer one can choose whether the first set is the empty or the full set (left and right of figure 8.2-A):

```
1       void first(ulong v=0)
2       {
3           ct_ = 0;
4           j_ = n_ - 1;
5           for (ulong j=0; j<n_; ++j)  x_[j] = v;
6       }
7
8       const ulong * data()  const  { return x_; }
9       ulong pos()  const { return j_; }
10      ulong current()  const  { return ct_; }
```

```
11
12      ulong next()
13      // Return position of change, return n with last subset
14      {
15          if ( ct_ == mct_ )  { return n_; }
16
17          ++ct_;
18          j_ = lowest_one_idx( ct_ );
19          x_[j_] ^= 1;
20
21          return  j_;
22      }
23
24      ulong prev()
25      // Return position of change, return n with first subset
26      {
27          if ( ct_ == 0 )  { return n_; }
28
29          j_ = lowest_one_idx( ct_ );
30          x_[j_] ^= 1;
31          --ct_;
32
33          return  j_;
34      }
35  };
```

About 130 million subsets are generated per second [FXT: comb/subset-gray-delta-demo.cc].

8.2.2 Generation as sets

The class [FXT: `class subset_gray` in comb/subset-gray.h] generates the subsets of $\{1, 2, \ldots, n\}$ as sets. The underlying idea is described in section 1.16.3 on page 43.

```
1   class subset_gray
2   // Subsets of the set {1,2,...,n} in minimal-change (Gray code) order.
3   {
4   public:
5       ulong *x_;  // data k-subset of {1,2,...,n} in x[1,...,k]
6       ulong n_;   // subsets of n-set
7       ulong k_;   // number of elements in subset
8
9   public:
10      subset_gray(ulong n)
11      {
12          n_ = n;
13          x_ = new ulong[n_+1];
14          x_[0] = 0;
15          first();
16      }
17
18      ~subset_gray()  { delete [] x_; }
19
20      ulong first()  { k_ = 0;  return k_; }
21      ulong last()  { x_[1] = 1;  k_ = 1;  return k_; }
22
23      const ulong * data() const { return x_+1; }
24      const ulong num() const { return k_; }
25
26  private:
27      ulong next_even()
28      {
29          if ( x_[k_]==n_ ) // remove n (from end):
30          {
31              --k_;
32          }
33          else  // append n:
34          {
35              ++k_;
36              x_[k_] = n_;
37          }
38          return  k_;
39      }
40
41      ulong next_odd()
42      {
```

```
43          if ( x_[k_]-1==x_[k_-1] ) // remove x[k]-1 (from position k-1):
44          {
45              x_[k_-1] = x_[k_];
46              --k_;
47          }
48          else // insert x[k]-1 as second last element:
49          {
50              x_[k_+1] = x_[k_];
51              --x_[k_];
52              ++k_;
53          }
54          return  k_;
55      }
56
57  public:
58      ulong next()
59      {
60          if ( 0==(k_&1 ) ) return next_even();
61          else              return next_odd();
62      }
63
64      ulong prev()
65      {
66          if ( 0==(k_&1 ) )  // k even
67          {
68              if ( 0==k_ )  return last();
69              return next_odd();
70          }
71          else  return next_even();
72      }
73  };
```

More than 160 million subsets are generated per second [FXT: comb/subset-gray-demo.cc]. The algorithm to generate the successor is given in [153].

8.2.3 Computing just the positions of change

The following algorithm computes only the locations of the changes, it is given in [43]. It can also be obtained as a specialization (for radix 2) of the loopless algorithm for computing a Gray code ordering of mixed radix numbers given section 9.2 on page 222 [FXT: `class ruler_func` in comb/ruler-func.h]:

```
1   class ruler_func
2   // Ruler function sequence: 0 1 0 2 0 1 0 3 0 1 0 2 0 1 0 4 0 1 0 2 0 1 ...
3   {
4   public:
5       ulong *f_;  // focus pointer
6       ulong n_;
7
8   public:
9       ruler_func(ulong n)
10      {
11          n_ = n;
12          f_ = new ulong[n+2];
13          first();
14      }
15
16      ~ruler_func()  { delete [] f_; }
17
18      void first()  { for (ulong k=0; k<n_+2; ++k)  f_[k] = k; }
19
20      ulong next()
21      {
22          const ulong j = f_[0];
23          //  if ( j==n_ )  { first(); return n_; }  // leave to user
24          f_[0] = 0;
25          const ulong j1 = j+1;
26          f_[j] = f_[j1];
27          f_[j1] = j1;
28          return j;
29      }
30  };
```

The sequence of positions is generated at a rate of about 290 million per second [FXT: comb/ruler-func-demo.cc].

8.3 Ordering with De Bruijn sequences

```
 0:    {0,  ,  ,  ,  }   #=1    {0}
 1:    { , 1,  ,  ,  }   #=1    {1}
 2:    { ,  , 2,  ,  }   #=1    {2}
 3:    { ,  ,  , 3,  }   #=1    {3}
 4:    {0,  ,  ,  , 4}   #=2    {0, 4}
 5:    {0, 1,  ,  ,  }   #=2    {0, 1}
 6:    { , 1, 2,  ,  }   #=2    {1, 2}
 7:    { ,  , 2, 3,  }   #=2    {2, 3}
 8:    {0,  ,  , 3, 4}   #=3    {0, 3, 4}
 9:    { , 1,  ,  , 4}   #=2    {1, 4}
10:    {0,  , 2,  ,  }   #=2    {0, 2}
11:    { , 1,  , 3,  }   #=2    {1, 3}
12:    { ,  , 2,  , 4}   #=2    {2, 4}
13:    {0,  ,  , 3,  }   #=2    {0, 3}
14:    {0, 1,  ,  , 4}   #=3    {0, 1, 4}
15:    {0, 1, 2,  ,  }   #=3    {0, 1, 2}
16:    { , 1, 2, 3,  }   #=3    {1, 2, 3}
17:    {0,  , 2, 3, 4}   #=4    {0, 2, 3, 4}
18:    { , 1,  , 3, 4}   #=3    {1, 3, 4}
19:    {0,  , 2,  , 4}   #=3    {0, 2, 4}
20:    {0, 1,  , 3,  }   #=3    {0, 1, 3}
21:    { , 1, 2,  , 4}   #=3    {1, 2, 4}
22:    {0,  , 2, 3,  }   #=3    {0, 2, 3}
23:    {0, 1,  , 3, 4}   #=4    {0, 1, 3, 4}
24:    {0, 1, 2,  , 4}   #=4    {0, 1, 2, 4}
25:    {0, 1, 2, 3,  }   #=4    {0, 1, 2, 3}
26:    {0, 1, 2, 3, 4}   #=5    {0, 1, 2, 3, 4}
27:    { , 1, 2, 3, 4}   #=4    {1, 2, 3, 4}
28:    { ,  , 2, 3, 4}   #=3    {2, 3, 4}
29:    { ,  ,  , 3, 4}   #=2    {3, 4}
30:    { ,  ,  ,  , 4}   #=1    {4}
31:    { ,  ,  ,  ,  }   #=0    {}
```

Figure 8.3-A: Subsets of a five element set in an order corresponding to a De Bruijn sequence.

A curious sequence of all subsets of a given set can be generated using a binary *De Bruijn sequence* that is a cyclical sequence of zeros and ones that contains each n-bit word once. In figure 8.3-A the empty places of the subsets are included to make the nice property apparent. The ordering has the *single track* property: each column in this (delta set) representation is a circular shift of the first column. The listing was created with the program [FXT: comb/subset-debruijn-demo.cc]. The underlying De Bruijn sequence is

```
0 0 0 0 0 1 0 0 0 1 1 0 0 1 0 1 0 0 1 1 1 0 1 0 1 1 0 1 1 1 1 1
```

(rotated left in the example so that the empty set appears at the end). Each subset is made from its predecessor by shifting it to the right and inserting the current element from the sequence.

The implementation [FXT: `class subset_debruijn` in comb/subset-debruijn.h] uses [FXT: `class debruijn` in comb/debruijn.h] (which in turn uses [FXT: `class necklace` in comb/necklace.h]). An algorithm for the generation of binary De Bruijn sequences is given in section 39.1 on page 857.

Successive subsets differ in many elements if the sequency (see section 1.17 on page 46) is big. Using the 'sequency-complemented' subsets (see end of section 1.17) we obtain an ordering where more elements change with small sequencies as shown in figure 8.3-B. This ordering corresponds to the complement-shift sequence of section 19.2.3 on page 379.

```
 0:    { ,  , 2,  , 4}   #=2    {2, 4}
 1:    {0, 1, 2,  , 4}   #=4    {0, 1, 2, 4}
 2:    {0,  ,  ,  , 4}   #=2    {0, 4}
 3:    {0,  , 2, 3, 4}   #=4    {0, 2, 3, 4}
 4:    { ,  , 2,  ,  }   #=1    {2}
 5:    { , 1, 2,  , 4}   #=3    {1, 2, 4}
 6:    {0, 1,  ,  , 4}   #=3    {0, 1, 4}
 7:    {0,  ,  , 3, 4}   #=3    {0, 3, 4}
 8:    { ,  , 2, 3,  }   #=2    {2, 3}
 9:    {0, 1, 2,  ,  }   #=3    {0, 1, 2}
10:    { ,  ,  ,  , 4}   #=1    {4}
11:    {0, 1, 2, 3, 4}   #=5    {0, 1, 2, 3, 4}
12:    {0,  ,  ,  ,  }   #=1    {0}
13:    { ,  , 2, 3, 4}   #=3    {2, 3, 4}
14:    { , 1, 2,  ,  }   #=2    {1, 2}
15:    { , 1,  ,  , 4}   #=2    {1, 4}
16:    {0, 1,  , 3, 4}   #=4    {0, 1, 3, 4}
17:    { ,  ,  , 3,  }   #=1    {3}
18:    {0, 1, 2, 3,  }   #=4    {0, 1, 2, 3}
19:    { ,  ,  ,  ,  }   #=0    {}
20:    { , 1, 2, 3, 4}   #=4    {1, 2, 3, 4}
21:    {0, 1,  ,  ,  }   #=2    {0, 1}
22:    { ,  ,  , 3, 4}   #=2    {3, 4}
23:    { , 1, 2, 3,  }   #=3    {1, 2, 3}
24:    { , 1,  ,  ,  }   #=1    {1}
25:    { , 1,  , 3, 4}   #=3    {1, 3, 4}
26:    { , 1,  , 3,  }   #=2    {1, 3}
27:    {0, 1,  , 3,  }   #=3    {0, 1, 3}
28:    {0,  ,  , 3,  }   #=2    {0, 3}
29:    {0,  , 2, 3,  }   #=3    {0, 2, 3}
30:    {0,  , 2,  ,  }   #=2    {0, 2}
31:    {0,  , 2,  , 4}   #=3    {0, 2, 4}
```

Figure 8.3-B: Subsets of a five element set in alternative order corresponding to a De Bruijn sequence. The elements 0, 2, and 4 are present exactly if they are not in figure 8.3-A.

8.4 Shifts-order for subsets

Figure 8.4-A shows an ordering (*shifts-order*) of the nonempty subsets of a 6-bit binary word where all linear shifts of a word appear in succession. The generation is done by a simple recursion [FXT: comb/shift-subsets-demo.cc]:

```
1   ulong n;  // number of bits
2   ulong N;  // 2**n
3
4   void A(ulong x)
5   {
6       if ( x>=N )  return;
7       visit(x);
8       A(2*x);
9       A(2*x+1);
10  }
```

The function `visit()` simply prints the binary expansion of its argument. The initial call is `A(1)`.

The transitions that are not shifts change just one bit if the following pair of functions is used for the recursion (*minimal-change shifts-order* shown in figure 8.4-B):

```
1   void F(ulong x)
2   {
3       if ( x>=N )  return;
4       visit(x);
5       F(2*x);
6       G(2*x+1);
7   }
8
9   void G(ulong x)
10  {
11      if ( x>=N )  return;
```

```
 1:  .....1   1    17:  1..111   4    33:  ....11   2    49:  ...111   3
 2:  ....1.   1    18:  ...1.1   2    34:  ...11.   2    50:  ..111.   3
 3:  ...1..   1    19:  ..1.1.   2    35:  ..11..   2    51:  .111..   3
 4:  ..1...   1    20:  .1.1..   2    36:  .11...   2    52:  111...   3
 5:  .1....   1    21:  1.1...   2    37:  11...1   2    53:  111..1   4
 6:  1.....   1    22:  1.1..1   3    38:  11...1   3    54:  .111.1   4
 7:  1....1   2    23:  .1.1.1   3    39:  .11..1   3    55:  111.1.   4
 8:  .1...1   2    24:  1.1.1.   3    40:  11..1.   3    56:  111.11   5
 9:  1...1.   2    25:  1.1.11   4    41:  11..11   4    57:  ..1111   4
10:  1...11   3    26:  ..1.11   3    42:  ..11.1   3    58:  .1111.   4
11:  ..1..1   2    27:  .1.11.   3    43:  .11.1.   3    59:  1111..   4
12:  .1..1.   2    28:  1.11..   3    44:  11.1..   3    60:  1111.1   5
13:  1..1..   2    29:  1.11.1   4    45:  11.1.1   4    61:  .11111   5
14:  1..1.1   3    30:  .1.111   4    46:  .11.11   4    62:  11111.   5
15:  .1..11   3    31:  1.111.   4    47:  11.11.   4    63:  111111   6
16:  1..11.   3    32:  1.1111   5    48:  11.111   5
```

Figure 8.4-A: Nonempty subsets of a 6-bit binary word where all linear shifts of a word appear in succession (shifts-order). All shifts are left shifts.

```
 1:  .....1   1    17:  ..1..1   2    33:  ...111   3    49:  .11.1.   3
 2:  ....1.   1    18:  ..1.11   3    34:  ..111.   3    50:  11.1..   3
 3:  ...1..   1    19:  .1.11.   3    35:  .111..   3    51:  11.1.1   4
 4:  ..1...   1    20:  1.11..   3    36:  111...   3    52:  11.111   5
 5:  .1....   1    21:  1.11.1   4    37:  111..1   4    53:  11.11.   4
 6:  1.....   1    22:  1.1111   5    38:  111.11   5    54:  .11.11   4
 7:  1....1   2    23:  1.111.   4    39:  111.1.   4    55:  .11..1   3
 8:  1...11   3    24:  .1.111   4    40:  .111.1   4    56:  11..1.   3
 9:  1...1.   2    25:  .1.1.1   3    41:  .11111   5    57:  11..11   4
10:  .1...1   2    26:  1.1.1.   3    42:  11111.   5    58:  11...1   3
11:  .1..11   3    27:  1.1.11   4    43:  111111   6    59:  11....   2
12:  1..11.   3    28:  1.1..1   3    44:  1111.1   5    60:  .11...   2
13:  1..111   4    29:  1.1...   2    45:  1111..   4    61:  ..11..   2
14:  1..1.1   3    30:  .1.1..   2    46:  .1111.   4    62:  ...11.   2
15:  1..1..   2    31:  ..1.1.   2    47:  ..1111   4    63:  ....11   2
16:  .1..1.   2    32:  ...1.1   2    48:  ..11.1   3
```

Figure 8.4-B: Nonempty subsets of a 6-bit binary word where all linear shifts of a word appear in succession and transitions that are not shifts switch just one bit (minimal-change shifts-order).

```
 1:  .......1   1    17:  ..1...1.   2    33:  1..1.1..   3    49:  ..1.1.1.   3
 2:  ......1.   1    18:  .1...1..   2    34:  1..1.1.1   4    50:  .1.1.1..   3
 3:  .....1..   1    19:  1...1...   2    35:  .....1.1   2    51:  1.1.1...   3
 4:  ....1...   1    20:  1...1..1   3    36:  ....1.1.   2    52:  1.1.1..1   4
 5:  ...1....   1    21:  .1...1.1   3    37:  ...1.1..   2    53:  .1.1.1.1   4
 6:  ..1.....   1    22:  1...1.1.   3    38:  ..1.1...   2    54:  1.1.1.1.   4
 7:  .1......   1    23:  ....1..1   2    39:  .1.1....   2
 8:  1.......   1    24:  ...1..1.   2    40:  1.1.....   2
 9:  1......1   2    25:  ..1..1..   2    41:  1.1....1   3
10:  .1.....1   2    26:  .1..1...   2    42:  .1.1...1   3
11:  1.....1.   2    27:  1..1....   2    43:  1.1...1.   3
12:  ..1....1   2    28:  1..1...1   3    44:  ..1.1..1   3
13:  .1....1.   2    29:  .1..1..1   3    45:  .1.1..1.   3
14:  1....1..   2    30:  1..1..1.   3    46:  1.1..1..   3
15:  1....1.1   3    31:  ..1..1.1   3    47:  1.1..1.1   4
16:  ...1...1   2    32:  .1..1.1.   3    48:  ...1.1.1   3
```

Figure 8.4-C: Nonzero Fibonacci words in an order where all shifts appear in succession.

```
12      F(2*x+1);
13      G(2*x);
14      visit(x);
15  }
```

The initial call is `F(1)`, the reversed order can be generated via `G(1)`.

We note that a simple variation can be used to generate the Fibonacci words in a shifts-order shown in figure 8.4-C. With transitions that are not shifts more than one bit is changed in general. The function used is [FXT: comb/shift-subsets-demo.cc]:

```
1   void B(ulong x)
2   {
3       if ( x>=N )  return;
4       visit(x);
5       B(2*x);
6       B(4*x+1);
7   }
```

An bit-level algorithm for combinations in shifts-order is given in section 1.26.3 on page 68.

8.5 k-subsets where k lies in a given range

We give algorithms for generating all *k-subsets* (subsets of an n-element set with k elements in each subset) where $k_{min} \le k \le k_{max}$. If $k_{min} = 0$ and $k_{max} = n$ we obtain all subsets, if $k_{min} = k_{max} = k$ we obtain combinations $\binom{n}{k}$.

8.5.1 Recursive algorithm

A recursive routine that generates all k-subsets where lies in a prescribed range is [FXT: `class ksubset_rec` in comb/ksubset-rec.h]. The routine can generate the subsets in 16 different orders. Figure 8.5-A shows the lexicographic orders, figure 8.5-B shows three Gray codes. The order numbers correspond to the second argument of the program [FXT: comb/ksubset-rec-demo.cc]. The constructor has just one argument, the number of elements of the set whose subsets shall be generated:

```
1   class ksubset_rec
2   // k-susbsets where kmin<=k<=kmax in various orders.
3   // Recursive CAT algorithm.
4   {
5   public:
6       long n_; // subsets of a n-element set
7       long kmin_, kmax_;  //  k-subsets where kmin<=k<=kma
8       long *rv_;  // record of visits in graph (list of elements in subset)
9       ulong ct_;   // count subsets
10      ulong rct_;  // count recursions (==work)
11      ulong rq_;   // condition that determines the order
12      ulong pq_;   // condition that determines the (printing) order
13      ulong nq_;   // whether to reverse order
14      // function to call with each combination:
15      void (*visit_)(const ksubset_rec &, long);
16
17  public:
18      ksubset_rec(ulong n)
19      {
20          n_ = n;
21          rv_ = new long[n_+1];
22          ++rv_;
23          rv_[-1] = -1UL;
24      }
25
26      ~ksubset_rec()
27      {
28          --rv_;
29          delete [] rv_;
30      }
```

```
          order #0:                                    order #8:
 0:   11....   ......   { 0, 1 }              111...   ......   { 0, 1, 2 }
 1:   111...   ..P...   { 0, 1, 2 }           11.1..   ..MP..   { 0, 1, 3 }
 2:   11.1..   ..MP..   { 0, 1, 3 }           11..1.   ...MP.   { 0, 1, 4 }
 3:   11..1.   ...MP.   { 0, 1, 4 }           11...1   ....MP   { 0, 1, 5 }
 4:   11...1   ....MP   { 0, 1, 5 }           11....   .....M   { 0, 1 }
 5:   1.1...   .MP..M   { 0, 2 }              1.11..   .MPP..   { 0, 2, 3 }
 6:   1.11..   ...P..   { 0, 2, 3 }           1.1.1.   ...MP.   { 0, 2, 4 }
 7:   1.1.1.   ...MP.   { 0, 2, 4 }           1.1..1   ....MP   { 0, 2, 5 }
 8:   1.1..1   ....MP   { 0, 2, 5 }           1.1...   .....M   { 0, 2 }
 9:   1..1..   ..MP.M   { 0, 3 }              1..11.   ..MPP.   { 0, 3, 4 }
10:   1..11.   ....P.   { 0, 3, 4 }           1..1.1   ....MP   { 0, 3, 5 }
11:   1..1.1   ....MP   { 0, 3, 5 }           1..1..   .....M   { 0, 3 }
12:   1...1.   ...MPM   { 0, 4 }              1...11   ...MPP   { 0, 4, 5 }
13:   1...11   .....P   { 0, 4, 5 }           1...1.   .....M   { 0, 4 }
14:   1....1   ....M.   { 0, 5 }              1....1   ....MP   { 0, 5 }
15:   .11...   MPP..M   { 1, 2 }              .111..   MPPP.M   { 1, 2, 3 }
16:   .111..   ...P..   { 1, 2, 3 }           .11.1.   ...MP.   { 1, 2, 4 }
17:   .11.1.   ...MP.   { 1, 2, 4 }           .11..1   ....MP   { 1, 2, 5 }
18:   .11..1   ....MP   { 1, 2, 5 }           .11...   .....M   { 1, 2 }
19:   .1.1..   ..MP.M   { 1, 3 }              .1.11.   ..MPP.   { 1, 3, 4 }
20:   .1.11.   ....P.   { 1, 3, 4 }           .1.1.1   ....MP   { 1, 3, 5 }
21:   .1.1.1   ....MP   { 1, 3, 5 }           .1.1..   .....M   { 1, 3 }
22:   .1..1.   ...MPM   { 1, 4 }              .1..11   ...MPP   { 1, 4, 5 }
23:   .1..11   .....P   { 1, 4, 5 }           .1..1.   .....M   { 1, 4 }
24:   .1...1   ....M.   { 1, 5 }              .1...1   ....MP   { 1, 5 }
25:   ..11..   .MPP.M   { 2, 3 }              ..111.   .MPPPM   { 2, 3, 4 }
26:   ..111.   ....P.   { 2, 3, 4 }           ..11.1   ....MP   { 2, 3, 5 }
27:   ..11.1   ....MP   { 2, 3, 5 }           ..11..   .....M   { 2, 3 }
28:   ..1.1.   ...MPM   { 2, 4 }              ..1.11   ...MPP   { 2, 4, 5 }
29:   ..1.11   .....P   { 2, 4, 5 }           ..1.1.   .....M   { 2, 4 }
30:   ..1..1   ....M.   { 2, 5 }              ..1..1   ....MP   { 2, 5 }
31:   ...11.   ..MPPM   { 3, 4 }              ...111   ..MPP.   { 3, 4, 5 }
32:   ...111   .....P   { 3, 4, 5 }           ...11.   .....M   { 3, 4 }
33:   ...1.1   ....M.   { 3, 5 }              ...1.1   ....MP   { 3, 5 }
34:   ....11   ...MP.   { 4, 5 }              ....11   ...MP.   { 4, 5 }
```

Figure 8.5-A: The k-subsets (where $2 \le k \le 3$) of a 6-element set. Lexicographic order for sets (left) and reversed lexicographic order for delta sets (right).

One has to supply the interval for k (variables `kmin` and `kmax`) and a function that will be called with each subset. The argument `rq` determines which of the sixteen different orderings is chosen, the order can be reversed with nonzero `nq`.

```
1       void generate(void (*visit)(const ksubset_rec &, long),
2                     long kmin, long kmax, ulong rq, ulong nq=0)
3       {
4           ct_ = 0;
5           rct_ = 0;
6
7           kmin_ = kmin;
8           kmax_ = kmax;
9           if ( kmin_ > kmax_ )  swap2(kmin_, kmax_);
10          if ( kmax_ > n_ )  kmax_ = n_;
11          if ( kmin_ > n_ )  kmin_ = n_;
12
13          visit_ = visit;
14          rq_ = rq % 4;
15          pq_ = (rq>>2) % 4;
16          nq_ = nq;
17          next_rec(0);
18      }
19
20  private:
21      void next_rec(long d);
22  };
```

The recursive routine itself is given in [FXT: comb/ksubset-rec.cc]:

```
1   void
```

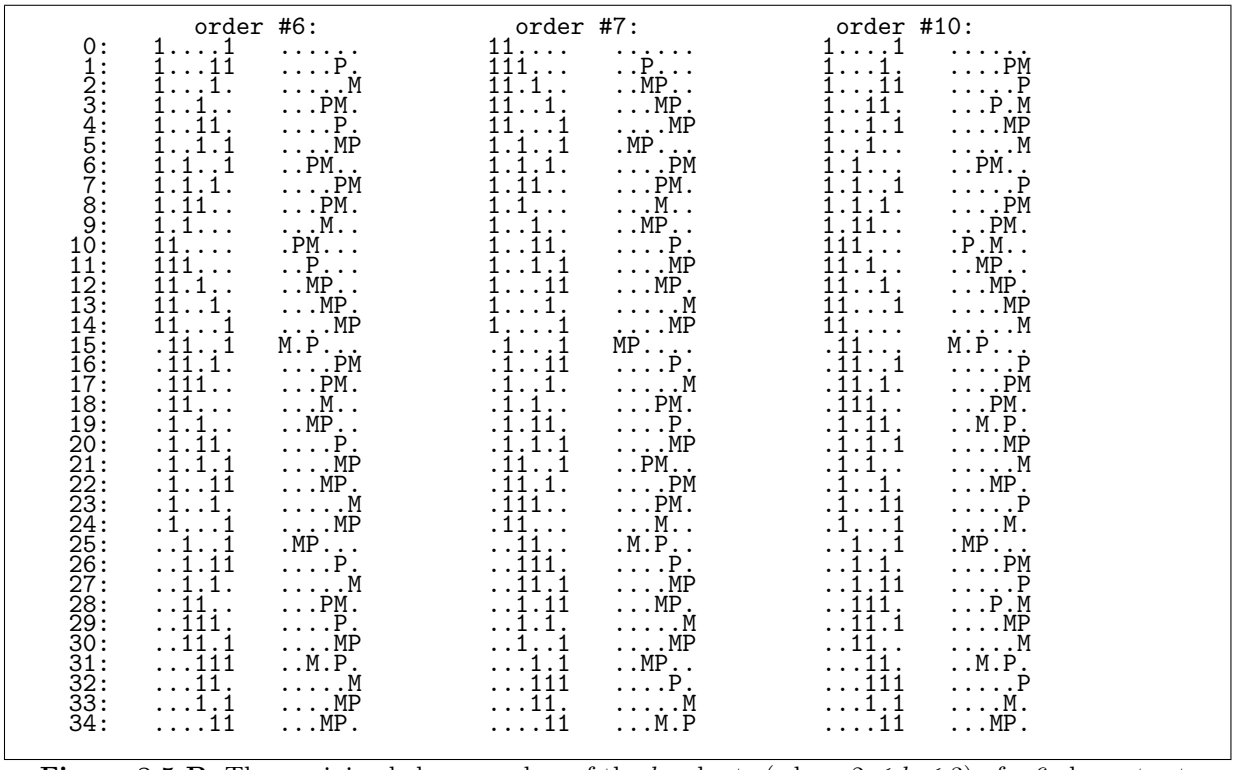

	order #6:		order #7:		order #10:	
0:	1....1	.....	11....	.....	1....1	.....
1:	1...11	....P.	111...	..P...	1...1.	....PM
2:	1...1.	.....M	11.1..	..MP..	1...11	.....P
3:	1..1..	...PM.	11..1.	...MP.	1..11.	...P.M
4:	1..11.	....P.	11...1	....MP	1..1.1	....MP
5:	1..1.1	....MP	1.1..1	.MP...	1..1..	.....M
6:	1.1..1	..PM..	1.1.1.	....PM	1.1...	..PM..
7:	1.1.1.	....PM	1.11..	...PM.	1.1..1	.....P
8:	1.11..	...PM.	1.1...	...M..	1.1.1.	....PM
9:	1.1...	...M..	1..1..	..MP..	1.11..	...PM.
10:	11....	.PM...	1..11.	....P.	111...	.P.M..
11:	111...	..P...	1..1.1	....MP	11.1..	..MP..
12:	11.1..	..MP..	1...11	...MP.	11..1.	...MP.
13:	11..1.	...MP.	1...1.	.....M	11...1	....MP
14:	11...1	....MP	1....1	....MP	11....	.....M
15:	.11..1	M.P...	.1...1	MP....	.11...	M.P...
16:	.11.1.	....PM	.1..11	....P.	.11..1	.....P
17:	.111..	...PM.	.1..1.	.....M	.11.1.	....PM
18:	.11...	...M..	.1.1..	...PM.	.111..	...PM.
19:	.1.1..	..MP..	.1.11.	....P.	.1.11.	..M.P.
20:	.1.11.	....P.	.1.1.1	....MP	.1.1.1	....MP
21:	.1.1.1	....MP	.11..1	..PM..	.1.1..	.....M
22:	.1..11	...MP.	.11.1.	....PM	.1..1.	...MP.
23:	.1..1.	.....M	.111..	...PM.	.1..11	.....P
24:	.1...1	....MP	.11...	...M..	.1...1	....M.
25:	..1..1	.MP...	..11..	.M.P..	..1..1	.MP...
26:	..1.11	....P.	..111.	....P.	..1.1.	....PM
27:	..1.1.	.....M	..11.1	....MP	..1.11	.....P
28:	..11..	...PM.	..1.11	...MP.	..111.	...P.M
29:	..111.	....P.	..1.1.	.....M	..11.1	....MP
30:	..11.1	....MP	..1..1	....MP	..11..	.....M
31:	...111	..M.P.	...1.1	..MP..	...11.	..M.P.
32:	...11.	.....M	...111	....P.	...111	.....P
33:	...1.1	....MP	...11.	.....M	...1.1	....M.
34:	11	...MP.	11	...M.P	11	...MP.

Figure 8.5-B: Three minimal-change orders of the k-subsets (where $2 \leq k \leq 3$) of a 6-element set.

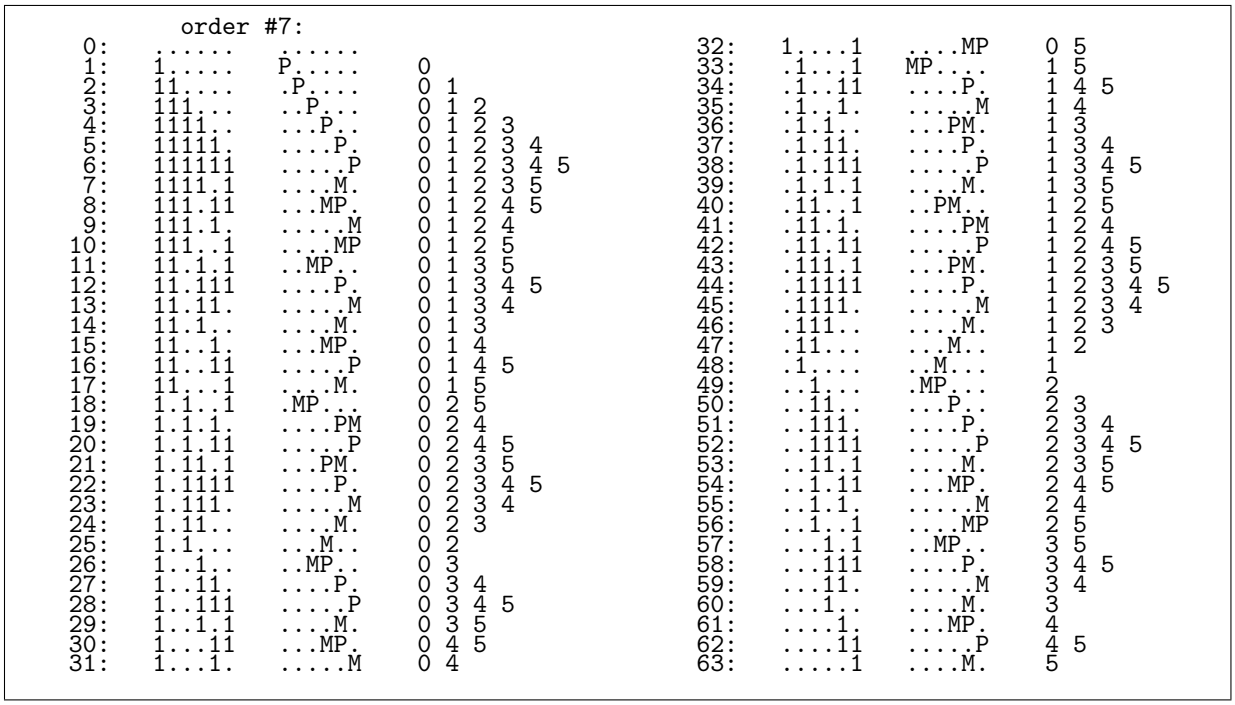

	order #7:														
0:	.....	.....							32:	1....1	....MP	0 5			
1:	1.....	P.....	0						33:	.1...1	MP....	1 5			
2:	11....	.P....	0 1						34:	.1..11	....P.	1 4 5			
3:	111...	..P...	0 1 2						35:	.1..1.	.....M	1 4			
4:	1111..	...P..	0 1 2 3						36:	.1.1..	...PM.	1 3			
5:	11111.	....P.	0 1 2 3 4						37:	.1.11.	....P.	1 3 4			
6:	111111	.....P	0 1 2 3 4 5						38:	.1.111	.....P	1 3 4 5			
7:	1111.1	....M.	0 1 2 3 5						39:	.1.1.1	....M.	1 3 5			
8:	111.11	...MP.	0 1 2 4 5						40:	.11..1	..PM..	1 2 5			
9:	111.1.	.....M	0 1 2 4						41:	.11.1.	....PM	1 2 4			
10:	111..1	....MP	0 1 2 5						42:	.11.11	.....P	1 2 4 5			
11:	11.1.1	..MP..	0 1 3 5						43:	.111.1	...PM.	1 2 3 5			
12:	11.111	....P.	0 1 3 4 5						44:	.11111	....P.	1 2 3 4 5			
13:	11.11.	.....M	0 1 3 4						45:	.1111.	.....M	1 2 3 4			
14:	11.1..	....M.	0 1 3						46:	.111..	....M.	1 2 3			
15:	11..1.	...MP.	0 1 4						47:	.11...	...M..	1 2			
16:	11..11	.....P	0 1 4 5						48:	.1....	..M...	1			
17:	11...1	....M.	0 1 5						49:	..1...	.MP...	2			
18:	1.1..1	.MP...	0 2 5						50:	..11..	...P..	2 3			
19:	1.1.1.	....PM	0 2 4						51:	..111.	....P.	2 3 4			
20:	1.1.11	.....P	0 2 4 5						52:	..1111	.....P	2 3 4 5			
21:	1.11.1	...PM.	0 2 3 5						53:	..11.1	....M.	2 3 5			
22:	1.1111	....P.	0 2 3 4 5						54:	..1.11	...MP.	2 4 5			
23:	1.111.	.....M	0 2 3 4						55:	..1.1.	.....M	2 4			
24:	1.11..	....M.	0 2 3						56:	..1..1	....MP	2 5			
25:	1.1...	...M..	0 2						57:	...1.1	..MP..	3 5			
26:	1..1..	..MP..	0 3						58:	...111	....P.	3 4 5			
27:	1..11.	....P.	0 3 4						59:	...11.	.....M	3 4			
28:	1..111	.....P	0 3 4 5						60:	...1..	....M.	3			
29:	1..1.1	....M.	0 3 5						61:	....1.	...MP.	4			
30:	1...11	...MP.	0 4 5						62:	11	.....P	4 5			
31:	1...1.	.....M	0 4						63:	1	....M.	5			

Figure 8.5-C: With $k_{min} = 0$ and order number seven at each transition either one element is added or removed, or one element moves to an adjacent position.

```
ksubset_rec::next_rec(long d)
{
    if ( d>kmax_ )  return;

    ++rct_;  // measure computational work
    long rv1 = rv_[d-1];  // left neighbor
    bool q;
    switch ( rq_ % 4 )
    {
    case 0:  q = 1;  break;
    case 1:  q = !(d&1);  break;
    case 2:  q = rv1&1;  break;
    case 3:  q = (d^rv1)&1;  break;
    }

    if ( nq_ )  q = !q;

    long x0 = rv1 + 1;
    long rx = n_ - (kmin_ - d);
    long x1 = min2( n_-1, rx );

#define PCOND(x) if ( (pq_==x) && (d>=kmin_) )  { visit_(*this, d);  ++ct_; }
    PCOND(0);
    if ( q )  // forward:
    {
        PCOND(1);
        for (long x=x0; x<=x1; ++x)  { rv_[d] = x;  next_rec(d+1); }
        PCOND(2);
    }
    else  // backward:
    {
        PCOND(2);
        for (long x=x1; x>=x0; --x)  { rv_[d] = x;  next_rec(d+1); }
        PCOND(1);
    }
    PCOND(3);
#undef PCOND
}
```

About 50 million subsets per second can be generated.

8.5.2 Iterative algorithm for a minimal-change order

The class [FXT: `class ksubset_gray` in comb/ksubset-gray.h] allows the generation of k-subsets of a set where k lies in a prescribed range:

```
class ksubset_gray
{
public:
    ulong n_;   // k-subsets of {1, 2, ..., n}
    ulong kmin_, kmax_;  // kmin <= k <= kmax
    ulong k_;   // k elements in current set
    ulong *S_;  // set in S[1,2,...,k] with elements \in {1,2,...,n}
    ulong j_;   // aux

public:
    ksubset_gray(ulong n, ulong kmin, ulong kmax)
    {
        n_ = (n>0 ? n : 1);
        // Must have 1<=kmin<=kmax<=n
        kmin_ = kmin;
        kmax_ = kmax;
        if ( kmax_ < kmin_ )  swap2(kmin_, kmax_);
        if ( kmin_==0 )  kmin_ = 1;

        S_ = new ulong[kmax_+1];
        S_[0] = 0;  // sentinel: != 1
        first();
    }

    ~ksubset_gray()  { delete [] S_; }
    const ulong *data()  const  { return S_+1; }
    const ulong num()  const  { return k_; }
```

```
           delta set  diff      set
       1:    ...11    .....    { 4, 5 }
       2:    ..11.    ..P.M    { 3, 4 }
       3:    ..111    ....P    { 3, 4, 5 }
       4:    ..1.1    ...M.    { 3, 5 }
       5:    .11..    .P..M    { 2, 3 }
       6:    .11.1    ....P    { 2, 3, 5 }
       7:    .1111    ...P.    { 2, 3, 4, 5 }
       8:    .111.    ....M    { 2, 3, 4 }
       9:    .1.1.    ..M..    { 2, 4 }
      10:    .1.11    ....P    { 2, 4, 5 }
      11:    .1..1    ...M.    { 2, 5 }
      12:    11...    P...M    { 1, 2 }
      13:    11..1    ....P    { 1, 2, 5 }
      14:    11.11    ...P.    { 1, 2, 4, 5 }
      15:    11.1.    ....M    { 1, 2, 4 }
      16:    1111.    ..P..    { 1, 2, 3, 4 }
      17:    111.1    ...MP    { 1, 2, 3, 5 }
      18:    111..    ....M    { 1, 2, 3 }
      19:    1.1..    .M...    { 1, 3 }
      20:    1.1.1    ....P    { 1, 3, 5 }
      21:    1.111    ...P.    { 1, 3, 4, 5 }
      22:    1.11.    ....M    { 1, 3, 4 }
      23:    1..1.    ..M..    { 1, 4 }
      24:    1..11    ....P    { 1, 4, 5 }
      25:    1...1    ...M.    { 1, 5 }
```

Figure 8.5-D: The (25) k-subsets where $2 \leq k \leq 4$ of a five-element set in a minimal-change order.

```
28
29      ulong last()
30      {
31          S_[1] = 1;  k_ = kmin_;
32          if ( kmin_==1 )  { j_ = 1; }
33          else
34          {
35              for (ulong i=2; i<=kmin_; ++i)  { S_[i] = n_ - kmin_ + i; }
36              j_ = 2;
37          }
38          return k_;
39      }
40
41
42      ulong first()
43      {
44          k_ = kmin_;
45          for (ulong i=1; i<=kmin_; ++i)  { S_[i] = n_ - kmin_ + i; }
46          j_ = 1;
47          return k_;
48      }
49
50      bool is_first()  const  { return ( S_[1] == n_ - kmin_ + 1 );  }
51
52      bool is_last()  const
53      {
54          if  ( S_[1] != 1  )   return 0;
55          if ( kmin_<=1 )  return (k_==1);
56          return  (S_[2]==n_-kmin_+2);
57      }
58    [--snip--]
```

The routines for computing the next or previous subset are adapted from a routine to compute the successor given in [153]. It is split in two auxiliary functions:

```
1   private:
2       void prev_even()
3       {
4           ulong &n=n_, &kmin=kmin_, &kmax=kmax_, &j=j_;
5           if ( S_[j-1] == S_[j]-1 )  // can touch sentinel
6           {
7               S_[j-1] = S_[j];
8               if ( j > kmin )
```

```
            {
                if ( S_[kmin] == n )  { j = j-2; }  else  { j = j-1; }
            }
            else
            {
                S_[j] = n - kmin + j;
                if ( S_[j-1]==S_[j]-1 )  { j = j-2; }
            }
        }
        else
        {
            S_[j] = S_[j] - 1;
            if ( j < kmax )
            {
                S_[j+1] = S_[j] + 1;
                if ( j >= kmin-1 )  { j = j+1; }  else  { j = j+2; }
            }
        }
    }

    void prev_odd()
    {
        ulong &n=n_, &kmin=kmin_, &kmax=kmax_, &j=j_;
        if ( S_[j] == n )  { j = j-1; }
        else
        {
            if ( j < kmax )
            {
                S_[j+1] = n;
                j = j+1;
            }
            else
            {
                S_[j] = S_[j]+1;
                if ( S_[kmin]==n )  { j = j-1; }
            }
        }
    }
[--snip--]
```

The `next()` and `prev()` functions use these routines, note that calls cannot not be mixed.

```
    ulong prev()
    {
        if ( is_first() )  { last(); return 0; }

        if ( j_&1 )  prev_odd();
        else         prev_even();

        if ( j_<kmin_ )  { k_ = kmin_; }  else  { k_ = j_; };

        return k_;
    }

    ulong next()
    {
        if ( is_last() )  { first();  return 0; }

        if ( j_&1 )  prev_even();
        else         prev_odd();

        if ( j_<kmin_ )  { k_ = kmin_; }  else  { k_ = j_; };

        return k_;
    }
[--snip--]
```

Usage of the class is shown in the program [FXT: comb/ksubset-gray-demo.cc], the k-subsets where $2 \le k \le 4$ in the order generated by the algorithm are shown in figure 8.5-D. About 80 million subsets per second can be generated with the routine `next()`, and 85 million with `prev()`.

	delta set	diff	set
1:	.1111	.....	{ 1, 2, 3, 4 }
2:	..111	.M...	{ 2, 3, 4 }
3:	1.111	P....	{ 0, 2, 3, 4 }
4:	11.11	.PM..	{ 0, 1, 3, 4 }
5:	.1.11	M....	{ 1, 3, 4 }
6:	...11	.M...	{ 3, 4 }
7:	1..11	P....	{ 0, 3, 4 }
8:	11..1	.P.M.	{ 0, 1, 4 }
9:	.1..1	M....	{ 1, 4 }
10:	1...1	PM...	{ 0, 4 }
11:	..1.1	M.P..	{ 2, 4 }
12:	1.1.1	P....	{ 0, 2, 4 }
13:	.11.1	MP...	{ 1, 2, 4 }
14:	111.1	P....	{ 0, 1, 2, 4 }
15:	1111.	...PM	{ 0, 1, 2, 3 }
16:	.111.	M....	{ 1, 2, 3 }
17:	..11.	.M...	{ 2, 3 }
18:	1.11.	P....	{ 0, 2, 3 }
19:	11.1.	.PM..	{ 0, 1, 3 }
20:	.1.1.	M....	{ 1, 3 }
21:	1..1.	PM...	{ 0, 3 }
22:	11...	.P.M.	{ 0, 1 }
23:	1.1..	.MP..	{ 0, 2 }
24:	.11..	MP...	{ 1, 2 }
25:	111..	P....	{ 0, 1, 2 }

	delta set	diff	set
1:	11	.....	{ 4, 5 }
2:	...1.1	...PM.	{ 3, 5 }
3:	.1...1	.P.M..	{ 1, 5 }
4:	1	.M....	{ 5 }
5:	1....1	P.....	{ 0, 5 }
6:	..1..1	M.P...	{ 2, 5 }
7:	..11..	...P.M	{ 2, 3 }
8:	.1.1..	.PM...	{ 1, 3 }
9:	...1..	.M....	{ 3 }
10:	1..1..	P.....	{ 0, 3 }
11:	11....	.P.M..	{ 0, 1 }
12:	.1....	M.....	{ 1 }
13:	1.....	PM....	{ 0 }
14:	..1...	M.P...	{ 2 }
15:	1.1...	P.....	{ 0, 2 }
16:	.11...	MP....	{ 1, 2 }
17:	.1..1.	..M.P.	{ 1, 4 }
18:	....1.	.M....	{ 4 }
19:	1...1.	P.....	{ 0, 4 }
20:	..1.1.	M.P...	{ 2, 4 }
21:	...11.	..MP..	{ 3, 4 }

Figure 8.5-E: The k-subsets where $2 \leq k \leq 4$ of 5 elements (left) and the sets where $1 \leq k \leq 2$ of 6 elements (right) in two-close orders.

8.5.3 A two-close order with homogenous moves

Orderings of the k-subsets with k in a given range that are *two-close* are shown in figure 8.5-E: one element is inserted or deleted or moves by at most two positions. The moves by two positions always cross a zero, the changes are *homogenous*. The list was produced with the program [FXT: comb/ksubset-twoclose-demo.cc] which uses [FXT: `class ksubset_twoclose` in comb/ksubset-twoclose.h]:

```
class ksubset_twoclose
// k-subsets (kmin<=k<=kmax) in a two-close order.
// Recursive algorithm.
{
public:
    ulong *rv_;  // record of visits in graph (delta set)
    ulong n_;    // subsets of the n-element set

    // function to call with each combination:
    void (*visit_)(const ksubset_twoclose &);
  [--snip--]

    void generate(void (*visit)(const ksubset_twoclose &),
                  ulong kmin, ulong kmax)
    {
        visit_ = visit;
        ulong kmax0 = n_ - kmin;
        next_rec(n_, kmax, kmax0, 0);
    }
```

The recursion is:

```
private:
    void next_rec(ulong d, ulong n1, ulong n0, bool q)
    // d:  remaining depth in recursion
    // n1: remaining ones to fill in
    // n0: remaining zeros to fill in
    // q:  direction in recursion
    {
        if ( 0==d )  { visit_(*this);  return; }

        --d;

```

```
12          if ( q )
13          {
14              if ( n0 )  { rv_[d]=0;  next_rec(d, n1-0, n0-1, d&1); }
15              if ( n1 )  { rv_[d]=1;  next_rec(d, n1-1, n0-0, q); }
16          }
17          else
18          {
19              if ( n1 )  { rv_[d]=1;  next_rec(d, n1-1, n0-0, q); }
20              if ( n0 )  { rv_[d]=0;  next_rec(d, n1-0, n0-1, d&1); }
21          }
22      }
23  };
```

About 50 million subsets per second can be generated. For $k_{min} = k_{max} =: k$ we obtain the enup order for combinations described in section 6.5.2 on page 186.

Chapter 9

Mixed radix numbers

The *mixed radix* representation $A = [a_0, a_1, a_2, \ldots, a_{n-1}]$ of a number x with respect to a radix vector $M = [m_0, m_1, m_2, \ldots, m_{n-1}]$ is given by the unique expression

$$x = \sum_{k=0}^{n-1} a_k \prod_{j=0}^{k-1} m_j \tag{9.0-1}$$

where $0 \le a_j < m_j$ (and $0 < x < \prod_{j=0}^{n-1} m_j$, so that n digits suffice). For $M = [r, r, r, \ldots, r]$ the relation reduces to the radix-r representation:

$$x = \sum_{k=0}^{n-1} a_k \, r^k \tag{9.0-2}$$

All 3-digit radix-4 numbers are shown in various orders in figure 9.0-A. Note that the least significant digit (a_0) is at the left side of each number (array representation).

9.1 Counting order

An implementation for mixed radix counting is [FXT: `class mixedradix_lex` in comb/mixedradix-lex.h]:

```
1   class mixedradix_lex
2   {
3   public:
4       ulong *a_;  // digits
5       ulong *m1_; // radix (minus one) for each digit
6       ulong n_;   // Number of digits
7       ulong j_;   // position of last change
8
9   public:
10      mixedradix_lex(const ulong *m, ulong n, ulong mm=0)
11      {
12          n_ = n;
13          a_ = new ulong[n_+1];
14          m1_ = new ulong[n_+1];
15          a_[n_] = 1;   // sentinel: !=0, and !=m1[n]
16          m1_[n_] = 0;  // sentinel
17          mixedradix_init(n_, mm, m, m1_);
18          first();
19      }
20    [--snip--]
```

The initialization routine is given in [FXT: comb/mixedradix-init.cc]:

	counting	Gray	modular Gray	gslex	endo	endo Gray
0:	[. . .]	[. . .]	[. . .]	[1 . .]	[. . .]	[. . .]
1:	[1 . .]	[1 . .]	[1 . .]	[2 . .]	[1 . .]	[1 . .]
2:	[2 . .]	[2 . .]	[2 . .]	[3 . .]	[3 . .]	[3 . .]
3:	[3 . .]	[3 . .]	[3 . .]	[1 1 .]	[2 . .]	[2 . .]
4:	[. 1 .]	[3 1 .]	[3 1 .]	[2 1 .]	[. 1 .]	[2 1 .]
5:	[1 1 .]	[2 1 .]	[. 1 .]	[3 1 .]	[1 1 .]	[3 1 .]
6:	[2 1 .]	[1 1 .]	[1 1 .]	[. 1 .]	[3 1 .]	[1 1 .]
7:	[3 1 .]	[. 1 .]	[2 1 .]	[1 2 .]	[2 1 .]	[. 1 .]
8:	[. 2 .]	[. 2 .]	[2 2 .]	[2 2 .]	[. 3 .]	[. 3 .]
9:	[1 2 .]	[1 2 .]	[3 2 .]	[3 2 .]	[1 3 .]	[1 3 .]
10:	[2 2 .]	[2 2 .]	[. 2 .]	[. 2 .]	[3 3 .]	[3 3 .]
11:	[3 2 .]	[3 2 .]	[1 2 .]	[1 3 .]	[2 3 .]	[2 3 .]
12:	[. 3 .]	[3 3 .]	[1 3 .]	[2 3 .]	[. 2 .]	[2 2 .]
13:	[1 3 .]	[2 3 .]	[2 3 .]	[3 3 .]	[1 2 .]	[3 2 .]
14:	[2 3 .]	[1 3 .]	[3 3 .]	[. 3 .]	[3 2 .]	[1 2 .]
15:	[3 3 .]	[. 3 .]	[. 3 .]	[1 . 1]	[2 2 .]	[. 2 .]
16:	[. . 1]	[. 3 1]	[. 3 1]	[2 . 1]	[. . 1]	[. 2 1]
17:	[1 . 1]	[1 3 1]	[1 3 1]	[3 . 1]	[1 . 1]	[1 2 1]
18:	[2 . 1]	[2 3 1]	[2 3 1]	[1 1 1]	[3 . 1]	[3 2 1]
19:	[3 . 1]	[3 3 1]	[3 3 1]	[2 1 1]	[2 . 1]	[2 2 1]
20:	[. 1 1]	[3 2 1]	[3 . 1]	[3 1 1]	[. 1 1]	[2 3 1]
21:	[1 1 1]	[2 2 1]	[. . 1]	[. 1 1]	[1 1 1]	[3 3 1]
22:	[2 1 1]	[1 2 1]	[1 . 1]	[1 2 1]	[3 1 1]	[1 3 1]
23:	[3 1 1]	[. 2 1]	[2 . 1]	[2 2 1]	[2 1 1]	[. 3 1]
24:	[. 2 1]	[. 1 1]	[2 1 1]	[3 2 1]	[. 3 1]	[. 1 1]
25:	[1 2 1]	[1 1 1]	[3 1 1]	[. 2 1]	[1 3 1]	[1 1 1]
26:	[2 2 1]	[2 1 1]	[. 1 1]	[1 3 1]	[3 3 1]	[3 1 1]
27:	[3 2 1]	[3 1 1]	[1 1 1]	[2 3 1]	[2 3 1]	[2 1 1]
28:	[. 3 1]	[3 . 1]	[1 2 1]	[3 3 1]	[. 2 1]	[2 . 1]
29:	[1 3 1]	[2 . 1]	[2 2 1]	[. 3 1]	[1 2 1]	[3 . 1]
30:	[2 3 1]	[1 . 1]	[3 2 1]	[. . 1]	[3 2 1]	[1 . 1]
31:	[3 3 1]	[. . 1]	[. 2 1]	[1 . 2]	[2 2 1]	[. . 1]
32:	[. . 2]	[. . 2]	[. 2 2]	[2 . 2]	[. . 3]	[. . 3]
33:	[1 . 2]	[1 . 2]	[1 2 2]	[3 . 2]	[1 . 3]	[1 . 3]
34:	[2 . 2]	[2 . 2]	[2 2 2]	[1 1 2]	[3 . 3]	[3 . 3]
35:	[3 . 2]	[3 . 2]	[3 2 2]	[2 1 2]	[2 . 3]	[2 . 3]
36:	[. 1 2]	[3 1 2]	[3 3 2]	[3 1 2]	[. 1 3]	[2 1 3]
37:	[1 1 2]	[2 1 2]	[. 3 2]	[. 1 2]	[1 1 3]	[3 1 3]
38:	[2 1 2]	[1 1 2]	[1 3 2]	[1 2 2]	[3 1 3]	[1 1 3]
39:	[3 1 2]	[. 1 2]	[2 3 2]	[2 2 2]	[2 1 3]	[. 1 3]
40:	[. 2 2]	[. 2 2]	[2 . 2]	[3 2 2]	[. 3 3]	[. 3 3]
41:	[1 2 2]	[1 2 2]	[3 . 2]	[. 2 2]	[1 3 3]	[1 3 3]
42:	[2 2 2]	[2 2 2]	[. . 2]	[1 3 2]	[3 3 3]	[3 3 3]
43:	[3 2 2]	[3 2 2]	[1 . 2]	[2 3 2]	[2 3 3]	[2 3 3]
44:	[. 3 2]	[3 3 2]	[1 1 2]	[3 3 2]	[. 2 3]	[2 2 3]
45:	[1 3 2]	[2 3 2]	[2 1 2]	[. 3 2]	[1 2 3]	[3 2 3]
46:	[2 3 2]	[1 3 2]	[3 1 2]	[. . 2]	[3 2 3]	[1 2 3]
47:	[3 3 2]	[. 3 2]	[. 1 2]	[1 . 3]	[2 2 3]	[. 2 3]
48:	[. . 3]	[. 3 3]	[. 1 3]	[2 . 3]	[. . 2]	[. 2 2]
49:	[1 . 3]	[1 3 3]	[1 1 3]	[3 . 3]	[1 . 2]	[1 2 2]
50:	[2 . 3]	[2 3 3]	[2 1 3]	[1 1 3]	[3 . 2]	[3 2 2]
51:	[3 . 3]	[3 3 3]	[3 1 3]	[2 1 3]	[2 . 2]	[2 2 2]
52:	[. 1 3]	[3 2 3]	[3 2 3]	[3 1 3]	[. 1 2]	[2 3 2]
53:	[1 1 3]	[2 2 3]	[. 2 3]	[. 1 3]	[1 1 2]	[3 3 2]
54:	[2 1 3]	[1 2 3]	[1 2 3]	[1 2 3]	[3 1 2]	[1 3 2]
55:	[3 1 3]	[. 2 3]	[2 2 3]	[2 2 3]	[2 1 2]	[. 3 2]
56:	[. 2 3]	[. 1 3]	[2 3 3]	[3 2 3]	[. 3 2]	[. 1 2]
57:	[1 2 3]	[1 1 3]	[3 3 3]	[. 2 3]	[1 3 2]	[1 1 2]
58:	[2 2 3]	[2 1 3]	[. 3 3]	[1 3 3]	[3 3 2]	[3 1 2]
59:	[3 2 3]	[3 1 3]	[1 3 3]	[2 3 3]	[2 3 2]	[2 1 2]
60:	[. 3 3]	[3 . 3]	[1 . 3]	[3 3 3]	[. 2 2]	[2 . 2]
61:	[1 3 3]	[2 . 3]	[2 . 3]	[. 3 3]	[1 2 2]	[3 . 2]
62:	[2 3 3]	[1 . 3]	[3 . 3]	[. . 3]	[3 2 2]	[1 . 2]
63:	[3 3 3]	[. . 3]	[. . 3]	[. . .]	[2 2 2]	[. . 2]

Figure 9.0-A: All 3-digit, radix-4 numbers in various orders (dots denote zeros): counting-, Gray-, modular Gray-, gslex-, endo-, and endo Gray order. The least significant digit is on the left of each word (array notation).

```
                M=[ 2 3 4 ]        M=[ 4 3 2 ]
          0:       [ . . . ]          [ . . . ]
          1:       [ 1 . . ]          [ 1 . . ]
          2:       [ . 1 . ]          [ 2 . . ]
          3:       [ 1 1 . ]          [ 3 . . ]
          4:       [ . 2 . ]          [ . 1 . ]
          5:       [ 1 2 . ]          [ 1 1 . ]
          6:       [ . . 1 ]          [ 2 1 . ]
          7:       [ 1 . 1 ]          [ 3 1 . ]
          8:       [ . 1 1 ]          [ . 2 . ]
          9:       [ 1 1 1 ]          [ 1 2 . ]
         10:       [ . 2 1 ]          [ 2 2 . ]
         11:       [ 1 2 1 ]          [ 3 2 . ]
         12:       [ . . 2 ]          [ . . 1 ]
         13:       [ 1 . 2 ]          [ 1 . 1 ]
         14:       [ . 1 2 ]          [ 2 . 1 ]
         15:       [ 1 1 2 ]          [ 3 . 1 ]
         16:       [ . 2 2 ]          [ . 1 1 ]
         17:       [ 1 2 2 ]          [ 1 1 1 ]
         18:       [ . . 3 ]          [ 2 1 1 ]
         19:       [ 1 . 3 ]          [ 3 1 1 ]
         20:       [ . 1 3 ]          [ . 2 1 ]
         21:       [ 1 1 3 ]          [ 1 2 1 ]
         22:       [ . 2 3 ]          [ 2 2 1 ]
         23:       [ 1 2 3 ]          [ 3 2 1 ]
```

Figure 9.1-A: Mixed radix numbers in counting order, dots denote zeros. The radix vectors are $M = [2, 3, 4]$ (rising factorial basis, left) and $M = [4, 3, 2]$ (falling factorial basis, right). The least significant digit is on the left of each word (array notation).

```
1   void
2   mixedradix_init(ulong n, ulong mm, const ulong *m, ulong *m1)
3   // Auxiliary function used to initialze vector of nines in mixed radix classes.
4   {
5       if ( m )  // all radices given
6       {
7           for (ulong k=0; k<n; ++k)  m1[k] = m[k] - 1;
8       }
9       else
10      {
11          if ( mm>1 )  // use mm as radix for all digits:
12              for (ulong k=0; k<n; ++k)  m1[k] = mm - 1;
13          else
14          {
15              if ( mm==0 )  // falling factorial basis
16                  for (ulong k=0; k<n; ++k)  m1[k] = n - k;
17              else // rising factorial basis
18                  for (ulong k=0; k<n; ++k)  m1[k] = k + 1;
19          }
20      }
21  }
```

Instead of the vector $M = [m_0, m_1, m_2, \ldots, m_{n-1}]$ the class uses the vector of 'nines', that is $M' = [m_0 - 1, m_1 - 1, m_2 - 1, \ldots, m_{n-1} - 1]$ (variable `m1_`). This modification leads to slightly faster generation. The first n-digit number is all-zero, the last is all-nines:

```
1       void first()
2       {
3           for (ulong k=0; k<n_; ++k)  a_[k] = 0;
4           j_ = n_;
5       }
6
7       void last()
8       {
9           for (ulong k=0; k<n_; ++k)  a_[k] = m1_[k];
10          j_ = n_;
11      }
12    [--snip--]
```

A number is incremented by setting all nines (digits a_j that are equal to $m_j - 1$) at the lower end to zero and incrementing the next digit:

```
1       bool next()  // increment
2       {
3           ulong j = 0;
4           while ( a_[j]==m1_[j] )  { a_[j]=0;  ++j; }  // can touch sentinels
5           j_ = j;
6
7           if ( j==n_ )  return false;  // current is last
8
9           ++a_[j];
10          return true;
11      }
12    [--snip--]
```

A number is decremented by setting all zero digits at the lower end to nine and decrementing the next digit:

```
1       bool prev()  // decrement
2       {
3           ulong j = 0;
4           while ( a_[j]==0 )  { a_[j]=m1_[j];  ++j; }  // can touch sentinels
5           j_ = j;
6
7           if ( j==n_ )  return false;  // current is first
8
9           --a_[j];
10          return true;
11      }
12    [--snip--]
```

Figure 9.1-A shows the 3-digit mixed radix numbers for basis vector $M = [2, 3, 4]$ (left) and $M = [4, 3, 2]$ (right). The listings where created with the program [FXT: comb/mixedradix-lex-demo.cc].

The routine `next()` generates between about 140 million (radix-2 numbers, $M = [2, 2, 2, \ldots, 2]$), 210 million (radix-3), and about 300 million (radix-8) numbers per second. Note that radix-2 leads to the slowest generation as the average carries are long compared to higher radices. The number of carries with incrementing is on average:

$$C \;=\; \frac{1}{m_0}\left(1+\frac{1}{m_1}\left(1+\frac{1}{m_2}\left(\ldots\right)\right)\right) \;=\; \sum_{k=0}^{n}\frac{1}{\prod_{j=0}^{k} m_j} \tag{9.1-1}$$

The number of digits changed on average equals $C + 1$. For $M = [r, r, r, \ldots, r]$ (and $n = \infty$) we obtain $C = \frac{1}{r-1}$. For the worst case ($r = 2$) we have $C = 1$, so two digits are changed on average.

9.2 Gray code order

Figure 9.2-A shows the 3-digit mixed radix numbers for radix vectors $M = [2, 3, 4]$ (left) and $M = [4, 3, 2]$ (right) in Gray code order. An constant amortized time (CAT) implementation for mixed radix numbers in a Gray code order is [FXT: `class mixedradix_gray` in comb/mixedradix-gray.h]:

```
1   class mixedradix_gray
2   {
3   public:
4       ulong *a_;  // mixed radix digits
5       ulong *m1_; // radices (minus one)
6       ulong *i_;  // direction
7       ulong n_;   // n_ digits
8       ulong j_;   // position of last change
9       int dm_;    // direction of last move
10
11  public:
12      mixedradix_gray(const ulong *m, ulong n, ulong mm=0)
13      {
14          n_ = n;
15          a_ = new ulong[n_+1];
16          a_[n] = -1UL;  // sentinel
17          i_ = new ulong[n_+1];
18          i_[n_] = 0;    // sentinel
```

```
          M=[ 2 3 4 ]     x     j  d          M=[ 4 3 2 ]     x     j  d
    0:       [ . . . ]     0                    [ . . . ]     0
    1:       [ 1 . . ]     1     0  1          [ 1 . . ]     1     0  1
    2:       [ 1 1 . ]     3     1  1          [ 2 . . ]     2     0  1
    3:       [ . 1 . ]     2     0 -1          [ 3 . . ]     3     0  1
    4:       [ . 2 . ]     4     1  1          [ 3 1 . ]     7     1  1
    5:       [ 1 2 . ]     5     0  1          [ 2 1 . ]     6     0 -1
    6:       [ 1 2 1 ]    11     2  1          [ 1 1 . ]     5     0 -1
    7:       [ . 2 1 ]    10     0 -1          [ . 1 . ]     4     0 -1
    8:       [ . 1 1 ]     8     1 -1          [ . 2 . ]     8     1  1
    9:       [ 1 1 1 ]     9     0  1          [ 1 2 . ]     9     0  1
   10:       [ 1 . 1 ]     7     1 -1          [ 2 2 . ]    10     0  1
   11:       [ . . 1 ]     6     0 -1          [ 3 2 . ]    11     0  1
   12:       [ . . 2 ]    12     2  1          [ 3 2 1 ]    23     2  1
   13:       [ 1 . 2 ]    13     0  1          [ 2 2 1 ]    22     0 -1
   14:       [ 1 1 2 ]    15     1  1          [ 1 2 1 ]    21     0 -1
   15:       [ . 1 2 ]    14     0 -1          [ . 2 1 ]    20     0 -1
   16:       [ . 2 2 ]    16     1  1          [ . 1 1 ]    16     1 -1
   17:       [ 1 2 2 ]    17     0  1          [ 1 1 1 ]    17     0  1
   18:       [ 1 2 3 ]    23     2  1          [ 2 1 1 ]    18     0  1
   19:       [ . 2 3 ]    22     0 -1          [ 3 1 1 ]    19     0  1
   20:       [ . 1 3 ]    20     1 -1          [ 3 . 1 ]    15     1 -1
   21:       [ 1 1 3 ]    21     0  1          [ 2 . 1 ]    14     0 -1
   22:       [ 1 . 3 ]    19     1 -1          [ 1 . 1 ]    13     0 -1
   23:       [ . . 3 ]    18     0 -1          [ . . 1 ]    12     0 -1
```

Figure 9.2-A: Mixed radix numbers in Gray code order, dots denote zeros. The radix vectors are $M = [2, 3, 4]$ (left) and $M = [4, 3, 2]$ (right). Columns '`x`' give the values, columns '`j`' and '`d`' give the position of last change and its direction, respectively.

```
19          m1_ = new ulong[n_+1];
20
21          mixedradix_init(n_, mm, m, m1_);
22
23          first();
24      }
25  [--snip--]
```

The array `i_[]` contains the 'directions' for each digits: it contains +1 or -1 if the computation of the successor will increase or decrease the corresponding digit. It has to be filled when the first or last number is computed:

```
1       void first()
2       {
3           for (ulong k=0; k<n_; ++k)  a_[k] = 0;
4           for (ulong k=0; k<n_; ++k)  i_[k] = +1;
5           j_ = n_;
6           dm_ = 0;
7       }
8
9       void last()
10      {
11          // find position of last even radix:
12          ulong z = 0;
13          for (ulong i=0; i<n_; ++i)  if ( m1_[i]&1 )  z = i;
14          while ( z<n_ )  // last even .. end:
15          {
16              a_[z] = m1_[z];
17              i_[z] = +1;
18              ++z;
19          }
20
21          j_ = 0;
22          dm_ = -1;
23      }
24  [--snip--]
```

A sentinel element (`i_[n]`=0) is used to optimize the computations of the successor and predecessor:

```
1       bool next()
2       {
3           ulong j = 0;
4           ulong ij;
```

```
5           while ( (ij=i_[j]) )  // can touch sentinel i[n]==0
6           {
7               ulong dj = a_[j] + ij;
8               if ( dj>m1_[j] )  // =^= if ( (dj>m1_[j]) || ((long)dj<0) )
9               {
10                  i_[j] = -ij;  // flip direction
11              }
12              else  // can update
13              {
14                  a_[j] = dj;  // update digit
15                  dm_ = ij;    // save for dir()
16                  j_ = j;      // save for pos()
17                  return true;
18              }
19
20              ++j;
21          }
22          return false;
23      }
24    [--snip--]
```

Note the if-clause, it is an optimized expression equivalent to the one given as comment. The following methods are often useful:

```
1       ulong pos()  const  { return j_; }  // position of last change
2       int dir()  const  { return dm_; }   // direction of last change
```

The routine for the computation of the predecessor is obtained by changing the statement `ulong dj = a_[j] + ij;` to `ulong dj = a_[j] - ij;`. About 120 million numbers per second for radix 2, and 245 million for radix 8 are generated [FXT: comb/mixedradix-gray-demo.cc].

A loopless algorithm for the computation of the successor taken from [170] is given in [FXT: comb/mixedradix-gray2.h]. It generates about 185 million numbers per second for radix 2, and 225 million for radix 8 [FXT: comb/mixedradix-gray2-demo.cc]. The crucial trick to make the algorithm loopless is the use of 'focus pointers':

```
1   class mixedradix_gray2
2   {
3   public:
4       ulong *a_;  // digits
5       ulong *m1_; // radix minus one ('nines')
6       ulong *f_;  // focus pointer
7       ulong *d_;  // direction
8       ulong n_;   // number of digits
9       ulong j_;   // position of last change
10      int dm_;    // direction of last move
11    [--snip--]
12      void first()
13      {
14          for (ulong k=0; k<n_; ++k)  a_[k] = 0;
15          for (ulong k=0; k<n_; ++k)  d_[k] = 1;
16          for (ulong k=0; k<=n_; ++k)  f_[k] = k;
17          dm_ = 0;
18          j_ = n_;
19      }
20
21      bool next()
22      {
23          const ulong j = f_[0];
24          f_[0] = 0;
25
26          if ( j>=n_ )  { first();  return false; }
27
28          const ulong dj = d_[j];
29          const ulong aj = a_[j] + dj;
30          a_[j] = aj;
31
32          dm_ = (int)dj;  // save for dir()
33          j_ = j;         // save for pos()
34
35          if ( aj+dj > m1_[j] )  // was last move?
36          {
37              d_[j] = -dj;      // change direction
38              f_[j] = f_[j+1];  // lookup next position
```

```
39              f_[j+1] = j + 1;
40          }
41
42          return true;
43      }
```

Modular Gray code order

```
              M=[ 2 3 4 ]     j                M=[ 4 3 2 ]     j
          0:    [ . . . ]                  0:    [ . . . ]
          1:    [ 1 . . ]     0            1:    [ 1 . . ]     0
          2:    [ 1 1 . ]     1            2:    [ 2 . . ]     0
          3:    [ . 1 . ]     0            3:    [ 3 . . ]     0
          4:    [ . 2 . ]     1            4:    [ 3 1 . ]     1
          5:    [ 1 2 . ]     0            5:    [ . 1 . ]     0
          6:    [ 1 2 1 ]     2            6:    [ 1 1 . ]     0
          7:    [ . 2 1 ]     0            7:    [ 2 1 . ]     0
          8:    [ . . 1 ]     1            8:    [ 2 2 . ]     1
          9:    [ 1 . 1 ]     0            9:    [ 3 2 . ]     0
         10:    [ 1 1 1 ]     1           10:    [ . 2 . ]     0
         11:    [ . 1 1 ]     0           11:    [ 1 2 . ]     0
         12:    [ . 1 2 ]     2           12:    [ 1 2 1 ]     2
         13:    [ 1 1 2 ]     0           13:    [ 2 2 1 ]     0
         14:    [ 1 2 2 ]     1           14:    [ 3 2 1 ]     0
         15:    [ . 2 2 ]     0           15:    [ . 2 1 ]     0
         16:    [ . . 2 ]     1           16:    [ . . 1 ]     1
         17:    [ 1 . 2 ]     0           17:    [ 1 . 1 ]     0
         18:    [ 1 . 3 ]     2           18:    [ 2 . 1 ]     0
         19:    [ . . 3 ]     0           19:    [ 3 . 1 ]     0
         20:    [ . 1 3 ]     1           20:    [ 3 1 1 ]     1
         21:    [ 1 1 3 ]     0           21:    [ . 1 1 ]     0
         22:    [ 1 2 3 ]     1           22:    [ 1 1 1 ]     0
         23:    [ . 2 3 ]     0           23:    [ 2 1 1 ]     0
```

Figure 9.2-B: Mixed radix numbers in Gray code order, dots denote zeros. The radix vectors are $M = [2, 3, 4]$ (left) and $M = [4, 3, 2]$ (right). The columns '`j`' give the position of last change.

Figure 9.2-B shows the 3-digit mixed radix numbers for radix vectors $M = [2, 3, 4]$ (left) and $M = [4, 3, 2]$ (right) in *modular Gray code* order. The transitions are either $k \to k + 1$ or, if k is maximal, $k \to 0$. The listing was created with the program [FXT: comb/mixedradix-modular-gray-demo.cc]. The loopless implementation [FXT: `class mixedradix_modular_gray` in comb/mixedradix-modular-gray.h] taken from [170] generates between about 135 million (radix 2) and 230 million (radix 16) numbers per second.

9.3 gslex order

The algorithm for the generation of subsets in lexicographic order given in section 8.1.2 on page 202 can be generalized for mixed radix numbers. Figure 9.3-A shows the 3-digit mixed radix numbers for basis vector $M = [2, 3, 4]$ (left) and $M = [4, 3, 2]$ (right). Note that zero is the last word in this order. For lack of a better name we call the order *gslex* (for generalized subset-lex) order. A routine for generating successive words in gslex order is implemented in [FXT: `class mixedradix_gslex` in comb/mixedradix-gslex.h]:

```
1   class mixedradix_gslex
2   {
3   public:
4       ulong n_;   // n-digit numbers
5       ulong *a_;  // digits
6       ulong *m1_; // m1[k] == radix-1 at position k
7
8   public:
9       mixedradix_gslex(ulong n, ulong mm, const ulong *m=0)
10      {
11          n_ = n;
```

```
              M=[ 2 3 4 ]     x                 M=[ 4 3 2 ]     x
        0:       [ 1 . . ]     1           0:       [ 1 . . ]     1
        1:       [ 1 1 . ]     3           1:       [ 2 . . ]     2
        2:       [ . 1 . ]     2           2:       [ 3 . . ]     3
        3:       [ 1 2 . ]     5           3:       [ 1 1 . ]     5
        4:       [ . 2 . ]     4           4:       [ 2 1 . ]     6
        5:       [ 1 . 1 ]     7           5:       [ 3 1 . ]     7
        6:       [ 1 1 1 ]     9           6:       [ . 1 . ]     4
        7:       [ . 1 1 ]     8           7:       [ 1 2 . ]     9
        8:       [ 1 2 1 ]    11           8:       [ 2 2 . ]    10
        9:       [ . 2 1 ]    10           9:       [ 3 2 . ]    11
       10:       [ . . 1 ]     6          10:       [ . 2 . ]     8
       11:       [ 1 . 2 ]    13          11:       [ 1 . 1 ]    13
       12:       [ 1 1 2 ]    15          12:       [ 2 . 1 ]    14
       13:       [ . 1 2 ]    14          13:       [ 3 . 1 ]    15
       14:       [ 1 2 2 ]    17          14:       [ 1 1 1 ]    17
       15:       [ . 2 2 ]    16          15:       [ 2 1 1 ]    18
       16:       [ . . 2 ]    12          16:       [ 3 1 1 ]    19
       17:       [ 1 . 3 ]    19          17:       [ . 1 1 ]    16
       18:       [ 1 1 3 ]    21          18:       [ 1 2 1 ]    21
       19:       [ . 1 3 ]    20          19:       [ 2 2 1 ]    22
       20:       [ 1 2 3 ]    23          20:       [ 3 2 1 ]    23
       21:       [ . 2 3 ]    22          21:       [ . 2 1 ]    20
       22:       [ . . 3 ]    18          22:       [ . . 1 ]    12
       23:       [ . . . ]     0          23:       [ . . . ]     0
```

Figure 9.3-A: Mixed radix numbers in gslex (generalized subset lex) order, dots denote zeros. The radix vectors are $M = [2, 3, 4]$ (left) and $M = [4, 3, 2]$ (right). Successive words differ in at most three positions. Columns 'x' give the values.

```
12          a_ = new ulong[n_ + 1];
13          a_[n_] = 1; // sentinel
14          m1_ = new ulong[n_];
15          mixedradix_init(n_, mm, m, m1_);
16          first();
17      }
18    [--snip--]
19      void first()
20      {
21          for (ulong k=0; k<n_; ++k)  a_[k] = 0;
22          a_[0] = 1;
23      }
24
25      void last()
26      {
27          for (ulong k=0; k<n_; ++k)  a_[k] = 0;
28      }
```

The method `next()` computes the successor:

```
1       bool next()
2       {
3           ulong e = 0;
4           while ( 0==a_[e] )  ++e;  // can touch sentinel
5
6           if ( e==n_ )  { first();  return false; }  // current is last
7
8           ulong ae = a_[e];
9           if ( ae != m1_[e] )  // easy case: simple increment
10          {
11              a_[0] = 1;
12              a_[e] = ae + 1;
13          }
14          else
15          {
16              a_[e] = 0;
17              if ( a_[e+1]==0 )  // can touch sentinel
18              {
19                  a_[0] = 1;
20                  ++a_[e+1];
21              }
22          }
23          return true;
```

```
24      }
```

The predecessor is computed by the method `prev()`:

```
1       bool prev()
2       {
3           ulong e = 0;
4           while ( 0==a_[e] )  ++e;  // can touch sentinel
5
6           if ( 0!=e )  // easy case: prepend nine
7           {
8               --e;
9               a_[e] = m1_[e];
10          }
11          else
12          {
13              ulong a0 = a_[0];
14              --a0;
15              a_[0] = a0;
16
17              if ( 0==a0 )
18              {
19                  do { ++e; }  while ( 0==a_[e] );  // can touch sentinel
20                  if ( e==n_ )  { last();  return false; }  // current is first
21                  ulong ae = a_[e];
22                  --ae;
23                  a_[e] = ae;
24                  if ( 0==ae )
25                  {
26                      --e;
27                      a_[e] = m1_[e];
28                  }
29              }
30          }
31          return true;
32      }
```

The algorithm is constant amortized time (CAT) and fast in practice. The worst performance occurs when all digits are radix 2, then about 123 million objects can be created per second. With radix 4 about 198 million, with radix 16 about 273 million objects per second are computed [FXT: comb/mixedradix-gslex-demo.cc].

Alternative gslex order

A variant of the gslex order is shown in figure 9.3-B. The ordering can be obtained from the gslex order by reversing the list, reversing the words, and replacing all nonzero digits d_i by $r_i - d_i$ where r_i is the radix at position i. The implementation is given in [FXT: `class mixedradix_gslex_alt` in comb/mixedradix-gslex-alt.h], the rate of generation is about the same as with gslex order [FXT: comb/mixedradix-gslex-alt-demo.cc].

9.4 endo order

The computation of the successor in mixed radix endo order (see section 6.5.1 on page 185) is very similar to the counting order described section 9.1 on page 219. The implementation [FXT: `class mixedradix_endo` in comb/mixedradix-endo.h] uses an additional array `le_[]` of the last nonzero elements in endo order. Its entries are 2 for $m > 1$, else 1:

```
1   class mixedradix_endo
2   {
3   public:
4       ulong *a_;  // digits, sentinel a[n]
5       ulong *m1_; // radix (minus one) for each digit
6       ulong *le_; // last positive digit in endo order, sentinel le[n]
7       ulong n_;   // Number of digits
8       ulong j_;   // position of last change
9
10      mixedradix_endo(const ulong *m, ulong n, ulong mm=0)
```

```
      M=[ 2 3 4 ]    x          M=[ 4 3 2 ]    x
 0:     [ . . . ]    0     0:     [ . . . ]    0
 1:     [ 1 . . ]    1     1:     [ 1 . . ]    1
 2:     [ 1 1 . ]    3     2:     [ 1 1 . ]    5
 3:     [ 1 1 1 ]    9     3:     [ 1 1 1 ]   17
 4:     [ 1 1 2 ]   15     4:     [ 1 2 . ]    9
 5:     [ 1 1 3 ]   21     5:     [ 1 2 1 ]   21
 6:     [ 1 2 . ]    5     6:     [ 1 . 1 ]   13
 7:     [ 1 2 1 ]   11     7:     [ 2 . . ]    2
 8:     [ 1 2 2 ]   17     8:     [ 2 1 . ]    6
 9:     [ 1 2 3 ]   23     9:     [ 2 1 1 ]   18
10:     [ 1 . 1 ]    7    10:     [ 2 2 . ]   10
11:     [ 1 . 2 ]   13    11:     [ 2 2 1 ]   22
12:     [ 1 . 3 ]   19    12:     [ 2 . 1 ]   14
13:     [ . 1 . ]    2    13:     [ 3 . . ]    3
14:     [ . 1 1 ]    8    14:     [ 3 1 . ]    7
15:     [ . 1 2 ]   14    15:     [ 3 1 1 ]   19
16:     [ . 1 3 ]   20    16:     [ 3 2 . ]   11
17:     [ . 2 . ]    4    17:     [ 3 2 1 ]   23
18:     [ . 2 1 ]   10    18:     [ 3 . 1 ]   15
19:     [ . 2 2 ]   16    19:     [ . 1 . ]    4
20:     [ . 2 3 ]   22    20:     [ . 1 1 ]   16
21:     [ . . 1 ]    6    21:     [ . 2 . ]    8
22:     [ . . 2 ]   12    22:     [ . 2 1 ]   20
23:     [ . . 3 ]   18    23:     [ . . 1 ]   12
```

Figure 9.3-B: Mixed radix numbers in alternative gslex (generalized subset lex) order, dots denote zeros. The radix vectors are $M = [2, 3, 4]$ (left) and $M = [4, 3, 2]$ (right). Successive words differ in at most three positions. Columns 'x' give the values.

```
      M=[ 5 6 ]    x                          x
 0:     [ . . ]    0    15:     [ . 5 ]   25
 1:     [ 1 . ]    1    16:     [ 1 5 ]   26
 2:     [ 3 . ]    3    17:     [ 3 5 ]   28
 3:     [ 4 . ]    4    18:     [ 4 5 ]   29
 4:     [ 2 . ]    2    19:     [ 2 5 ]   27
 5:     [ . 1 ]    5    20:     [ . 4 ]   20
 6:     [ 1 1 ]    6    21:     [ 1 4 ]   21
 7:     [ 3 1 ]    8    22:     [ 3 4 ]   23
 8:     [ 4 1 ]    9    23:     [ 4 4 ]   24
 9:     [ 2 1 ]    7    24:     [ 2 4 ]   22
10:     [ . 3 ]   15    25:     [ . 2 ]   10
11:     [ 1 3 ]   16    26:     [ 1 2 ]   11
12:     [ 3 3 ]   18    27:     [ 3 2 ]   13
13:     [ 4 3 ]   19    28:     [ 4 2 ]   14
14:     [ 2 3 ]   17    29:     [ 2 2 ]   12
```

Figure 9.4-A: Mixed radix numbers in endo order, dots denote zeros. The radix vector is $M = [5, 6]$. Columns 'x' give the values.

```
11      {
12          n_ = n;
13          a_ = new ulong[n_+1];
14          a_[n_] = 1;  // sentinel:  != 0
15          m1_ = new ulong[n_];
16
17          mixedradix_init(n_, mm, m, m1_);
18
19          le_ = new ulong[n_+1];
20          le_[n_] = 0;  // sentinel:  != a[n]
21          for (ulong k=0; k<n_; ++k)  le_[k] = 2 - (m1_[k]==1);
22
23          first();
24      }
25    [--snip--]
```

The first number is all zero, the last can be read form the array `le_[]`:

```
1       void first()
2       {
3           for (ulong k=0; k<n_; ++k)  a_[k] = 0;
4           j_ = n_;
5       }
6
7       void last()
8       {
9           for (ulong k=0; k<n_; ++k)  a_[k] = le_[k];
10          j_ = n_;
11      }
12    [--snip--]
```

In the computation of the successor the function `next_endo()` is used instead of a simple increment:

```
1       bool next()
2       {
3           bool ret = false;
4           ulong j = 0;
5
6           while ( a_[j]==le_[j] )  { a_[j]=0; ++j; }  // can touch sentinel
7           if ( j<n_ ) // only if no overflow
8           {
9               a_[j] = next_endo(a_[j], m1_[j]);  // increment
10              ret = true;
11          }
12
13          j_ = j;
14          return ret;
15      }
16
17      bool prev()
18      {
19          bool ret = false;
20          ulong j = 0;
21
22          while ( a_[j]==0 )  { a_[j]=le_[j]; ++j; }  // can touch sentinel
23          if ( j<n_ ) // only if no overflow
24          {
25              a_[j] = prev_endo(a_[j], m1_[j]);  // decrement
26              ret = true;
27          }
28
29          j_ = j;
30          return ret;
31      }
32    [--snip--]
```

The function `next()` generates between about 115 million (radix 2) and 180 million (radix 16) numbers per second. The listing in figure 9.4-A was created with the program [FXT: comb/mixedradix-endo-demo.cc].

```
       M=[ 5 6 ]     x     j    d                          x     j    d
  0:     [ . . ]     0                 15:    [ 2 5 ]    27     1    1
  1:     [ 1 . ]     1     0    1      16:    [ 4 5 ]    29     0   -1
  2:     [ 3 . ]     3     0    1      17:    [ 3 5 ]    28     0   -1
  3:     [ 4 . ]     4     0    1      18:    [ 1 5 ]    26     0   -1
  4:     [ 2 . ]     2     0    1      19:    [ . 5 ]    25     0   -1
  5:     [ 2 1 ]     7     1    1      20:    [ . 4 ]    20     1    1
  6:     [ 4 1 ]     9     0   -1      21:    [ 1 4 ]    21     0    1
  7:     [ 3 1 ]     8     0   -1      22:    [ 3 4 ]    23     0    1
  8:     [ 1 1 ]     6     0   -1      23:    [ 4 4 ]    24     0    1
  9:     [ . 1 ]     5     0   -1      24:    [ 2 4 ]    22     0    1
 10:     [ . 3 ]    15     1    1      25:    [ 2 2 ]    12     1    1
 11:     [ 1 3 ]    16     0    1      26:    [ 4 2 ]    14     0   -1
 12:     [ 3 3 ]    18     0    1      27:    [ 3 2 ]    13     0   -1
 13:     [ 4 3 ]    19     0    1      28:    [ 1 2 ]    11     0   -1
 14:     [ 2 3 ]    17     0    1      29:    [ . 2 ]    10     0   -1
```

Figure 9.5-A: Mixed radix numbers in endo Gray code, dots denote zeros. The radix vector is $M = [4, 5]$. Columns '`x`' give the values, columns '`j`' and '`d`' give the position of last change and its direction, respectively.

9.5 Gray code for endo order

A Gray code for mixed radix numbers in endo order can be obtained by a modification of the CAT algorithm for the Gray code described in section 9.2 on page 222. In the computation of the last number, the last digit have to be set to the last endo digit [FXT: `class mixedradix_endo_gray` in comb/mixedradix-endo-gray.h]:

```
1   class mixedradix_endo_gray
2   {
3   public:
4       ulong *a_;  // mixed radix digits
5       ulong *m1_; // radices (minus one)
6       ulong *i_;  // direction
7       ulong *le_; // last positive digit in endo order
8       ulong n_;   // n_ digits
9       ulong j_;   // position of last change
10      int dm_;    // direction of last move
11
12    [--snip--]
13      void first()
14      {
15          for (ulong k=0; k<n_; ++k)  a_[k] = 0;
16          for (ulong k=0; k<n_; ++k)  i_[k] = +1;
17          j_ = n_;
18          dm_ = 0;
19      }
20
21      void last()
22      {
23          for (ulong k=0; k<n_; ++k)  a_[k] = 0;
24          for (ulong k=0; k<n_; ++k)  i_[k] = -1UL;
25
26          // find position of last even radix:
27          ulong z = 0;
28          for (ulong i=0; i<n_; ++i)  if ( m1_[i]&1 )  z = i;
29          while ( z<n_ )  // last even .. end:
30          {
31              a_[z] = m1_[z];
32              i_[z] = +1;
33              ++z;
34          }
35
36          j_ = 0;
37          dm_ = -1;
38      }
39    [--snip--]
```

The successor is computed as follows:

```
1       bool next()
```

```
2       {
3           ulong j = 0;
4           ulong ij;
5           while ( (ij=i_[j]) )  // can touch sentinel i[n]==0
6           {
7               ulong dj;
8               bool ovq;  // overflow?
9               if ( ij == 1 )
10              {
11                  dj = next_endo(a_[j], m1_[j]);
12                  ovq = (dj==0);
13              }
14              else
15              {
16                  ovq = (a_[j]==0);
17                  dj = prev_endo(a_[j], m1_[j]);
18              }
19
20              if ( ovq )  i_[j] = -ij;
21              else
22              {
23                  a_[j] = dj;
24                  dm_ = ij;
25                  j_ = j;
26                  return true;
27              }
28
29              ++j;
30          }
31          return false;
32      }
33    [--snip--]
```

The routine for computation of the predecessor is obtained by changing the condition `if ( ij == 1 )` to `if ( ij != 1 )`. About 65 million (radix 2) and 110 million (radix 16) numbers per second are generated. The listing in figure 9.5-A was created with the program [FXT: comb/mixedradix-endo-gray-demo.cc].

Chapter 10

Permutations

In this section algorithms for the generation of all permutations are presented. These are typically useful in situations where an exhaustive search over all permutations is needed. Some algorithms use mixed radix numbers with factorial base, see section 10.3 on page 236.

Algorithms for application, inversion and composition of permutations and the generation of random permutations are given in chapter 2 on page 93. The sign (parity) of a permutation is defined in section 2.11.5 on page 116.

A important optimization technique is to use arrays instead of pointers. One would change the pointer declarations to array declarations in the corresponding class as follows:

```
//ulong *p_;   // permutation data (pointer version)
ulong p_[32];  // permutation data (array version)
```

One also needs to disables the statements to allocate and free memory with the pointers. Here we assume that nobody would attempt to compute all permutations of 31 or more elements ($31! \approx 8.22 \cdot 10^{33}$, taking about $1.3 \cdot 10^{18}$ years to finish). To use arrays uncomment (in most implementations) a line like

```
#define PERM_REV2_FIXARRAYS // use arrays instead of pointers (speedup)
```

near the top of the header file. Whether the use of arrays tends to give a speedup is noted in the comment, as above.

10.1 Lexicographic order

When generated in lexicographic order the permutations appear as if (read as numbers and) sorted numerically in ascending order, see figure 10.1-A. The first half of the inverse permutations are the reversed inverse permutations in the second half: the position of zero in the first half of the inverse permutations lies in the first half of each permutation, so their reversal gives the second half. Write I for the operator that inverts a permutation, C for the complement, and R for reversal. Then we have

$$C \;=\; I\,R\,I \tag{10.1-1}$$

and thereby the first half of the permutations are the complements of the permutations in the second half. An implementation of an iterative algorithm is [FXT: `class perm_lex` in comb/perm-lex.h].

```
1  class perm_lex
2  {
3  public:
4      ulong *p_;  // permutation in 0, 1, ..., n-1, sentinel at [-1]
5      ulong n_;   // number of elements to permute
6
7  public:
8      perm_lex(ulong n)
```

	permutation	inv. perm.	compl. inv. perm.	reversed perm.
0:	[. 1 2 3]	[. 1 2 3]	[3 2 1 .]	[3 2 1 .]
1:	[. 1 3 2]	[. 1 3 2]	[3 2 . 1]	[2 3 1 .]
2:	[. 2 1 3]	[. 2 1 3]	[3 1 2 .]	[3 1 2 .]
3:	[. 2 3 1]	[. 3 1 2]	[3 . 2 1]	[1 3 2 .]
4:	[. 3 1 2]	[. 2 3 1]	[3 1 . 2]	[2 1 3 .]
5:	[. 3 2 1]	[. 3 2 1]	[3 . 1 2]	[1 2 3 .]
6:	[1 . 2 3]	[1 . 2 3]	[2 3 1 .]	[3 2 . 1]
7:	[1 . 3 2]	[1 . 3 2]	[2 3 . 1]	[2 3 . 1]
8:	[1 2 . 3]	[2 . 1 3]	[1 3 2 .]	[3 . 2 1]
9:	[1 2 3 .]	[3 . 1 2]	[. 3 2 1]	[. 3 2 1]
10:	[1 3 . 2]	[2 . 3 1]	[1 3 . 2]	[2 . 3 1]
11:	[1 3 2 .]	[3 . 2 1]	[. 3 1 2]	[. 2 3 1]
12:	[2 . 1 3]	[1 2 . 3]	[2 1 3 .]	[3 1 . 2]
13:	[2 . 3 1]	[1 3 . 2]	[2 . 3 1]	[1 3 . 2]
14:	[2 1 . 3]	[2 1 . 3]	[1 2 3 .]	[3 . 1 2]
15:	[2 1 3 .]	[3 1 . 2]	[. 2 3 1]	[. 3 1 2]
16:	[2 3 . 1]	[2 3 . 1]	[1 . 3 2]	[1 . 3 2]
17:	[2 3 1 .]	[3 2 . 1]	[. 1 3 2]	[. 1 3 2]
18:	[3 . 1 2]	[1 2 3 .]	[2 1 . 3]	[2 1 . 3]
19:	[3 . 2 1]	[1 3 2 .]	[2 . 1 3]	[1 2 . 3]
20:	[3 1 . 2]	[2 1 3 .]	[1 2 . 3]	[2 . 1 3]
21:	[3 1 2 .]	[3 1 2 .]	[. 2 1 3]	[. 2 1 3]
22:	[3 2 . 1]	[2 3 1 .]	[1 . 2 3]	[1 . 2 3]
23:	[3 2 1 .]	[3 2 1 .]	[. 1 2 3]	[. 1 2 3]

Figure 10.1-A: All permutations of 4 elements in lexicographic order, their inverses, the complements of the inverses, and the reversed permutations. Dots denote zeros.

```
9       {
10          n_ = n;
11          p_ = new ulong[n_+1];
12          p_[0] = 0;  // sentinel
13          ++p_;
14          first();
15      }
16
17      ~perm_lex() { --p_;  delete [] p_; }
18
19      void first()  { for (ulong i=0; i<n_; i++)  p_[i] = i; }
20
21      const ulong *data()  const  { return p; }
22    [--snip--]
```

The only nontrivial part is the `next()`-method that computes the next permutation with each call. The routine `perm_lex::next()` is based on code by Glenn Rhoads

```
1       bool next()
2       {
3           // find for rightmost pair with p_[i] < p_[i+1]:
4           const ulong n1 = n_ - 1;
5           ulong i = n1;
6           do  { --i; }  while ( p_[i] > p_[i+1] );
7           if ( (long)i<0 )  return false;  // last sequence is falling seq.
8
9           // find rightmost element p[j] smaller than p[i]:
10          ulong j = n1;
11          while ( p_[i] > p_[j] )  { --j; }
12
13          swap2(p_[i], p_[j]);
14
15          // Here the elements p[i+1], ..., p[n-1] are a falling sequence.
16          // Reverse order to the right:
17          ulong r = n1;
18          ulong s = i + 1;
19          while ( r > s )  { swap2(p_[r], p_[s]);  --r;  ++s; }
20
21          return true;
22      }
```

The routine generates about 113 million permutations per second. Using the class is no black magic [FXT: comb/perm-lex-demo.cc]:

```
ulong n = 4;
perm_lex P(n);
do
{
    // visit permutation
}
while ( P.next() );
```

A slightly faster algorithm is obtained by making the changes with the update operation for the co-lexicographic order (section 10.2) on the right end of the permutations [FXT: comb/perm-lex2.h]. When arrays are used instead of pointers the rate is about 133 million per second [FXT: comb/perm-lex2-demo.cc]. With pointers the rate is about 115 million per second.

10.2 Co-lexicographic order

	permutation	rfact	inv. perm.
0:	[3 2 1 .]	[. . .]	[3 2 1 .]
1:	[2 3 1 .]	[1 . .]	[3 2 . 1]
2:	[3 1 2 .]	[. 1 .]	[3 1 2 .]
3:	[1 3 2 .]	[1 1 .]	[3 . 2 1]
4:	[2 1 3 .]	[. 2 .]	[3 1 . 2]
5:	[1 2 3 .]	[1 2 .]	[3 . 1 2]
6:	[3 2 . 1]	[. . 1]	[2 3 1 .]
7:	[2 3 . 1]	[1 . 1]	[2 3 . 1]
8:	[3 . 2 1]	[. 1 1]	[1 3 2 .]
9:	[. 3 2 1]	[1 1 1]	[. 3 2 1]
10:	[2 . 3 1]	[. 2 1]	[1 3 . 2]
11:	[. 2 3 1]	[1 2 1]	[. 3 1 2]
12:	[3 1 . 2]	[. . 2]	[2 1 3 .]
13:	[1 3 . 2]	[1 . 2]	[2 . 3 1]
14:	[3 . 1 2]	[. 1 2]	[1 2 3 .]
15:	[. 3 1 2]	[1 1 2]	[. 2 3 1]
16:	[1 . 3 2]	[. 2 2]	[1 . 3 2]
17:	[. 1 3 2]	[1 2 2]	[. 1 3 2]
18:	[2 1 . 3]	[. . 3]	[2 1 . 3]
19:	[1 2 . 3]	[1 . 3]	[2 . 1 3]
20:	[2 . 1 3]	[. 1 3]	[1 2 . 3]
21:	[. 2 1 3]	[1 1 3]	[. 2 1 3]
22:	[1 . 2 3]	[. 2 3]	[1 . 2 3]
23:	[. 1 2 3]	[1 2 3]	[. 1 2 3]

Figure 10.2-A: The permutations of 4 elements in co-lexicographic order. Dots denote zeros.

Figure 10.2-A shows the permutations of 4 elements in co-lexicographic (colex) order. An algorithm for the generation is implemented in [FXT: `class perm_colex` in comb/perm-colex.h]:

```
1   class perm_colex
2   {
3   public:
4       ulong *d_;  // mixed radix digits with radix = [2, 3, 4, ...]
5       ulong *x_;  // permutation
6       ulong n_;   // permutations of n elements
7
8   public:
9       perm_colex(ulong n)
10      // Must have n>=2
11      {
12          n_ = n;
13          d_ = new ulong[n_];
14          d_[n-1] = 0;  // sentinel
15          x_ = new ulong[n_];
16          first();
17      }
18    [--snip--]
19
20      void first()
21      {
22          for (ulong k=0; k<n_; ++k)  x_[k] = n_-1-k;
```

```
23          for (ulong k=0; k<n_-1; ++k)  d_[k] = 0;
24      }
25
```

The update process uses the falling factorial numbers. Let j be the position where the digit is incremented, and d the value before the increment. The update

```
  permutation       ffact
                          v-- increment at j=3
[ 0 3 4 5 2 1 ]   [ 1 2 3 1 1 ]  <--= digit before increment is d=1
[ 5 4 2 0 3 1 ]   [ . . . 2 1 ]
```

is done in three steps:

```
[ 0 3 4 5 2 1 ]   [ 1 2 3 1 1 ]
[ 0 2 4 5 3 1 ]   [ 1 2 3 2 1 ]  <--= swap positions d=1 and j+1=4
[ 5 4 2 0 3 1 ]   [ . . . 2 1 ]  <--= reverse range 0...j
```

```
1       bool next()
2       {
3           if ( d_[0]==0 ) // easy case
4           {
5               d_[0] = 1;
6               swap2(x_[0], x_[1]);
7               return true;
8           }
9           else
10          {
11              d_[0] = 0;
12              ulong j = 1;
13              ulong m1 = 2;  // nine in falling factorial base
14              while ( d_[j]==m1 )
15              {
16                  d_[j] = 0;
17                  ++m1;
18                  ++j;
19              }
20
21              if ( j==n_-1 )  return false;  // current permutation is last
22
23              const ulong dj = d_[j];
24              d_[j] = dj + 1;
25
26              swap2( x_[dj], x_[j+1] );  // swap positions dj and j+1
27
28              { // reverse range [0...j]:
29                  ulong a = 0,  b = j;
30                  do
31                  {
32                      swap2(x_[a], x_[b]);
33                      ++a;
34                      --b;
35                  }
36                  while ( a<b );
37              }
38
39              return true;
40          }
41      }
42  }
```

About 194 million permutations per second can be generated [FXT: comb/perm-colex-demo.cc]. With arrays instead of pointers the rate is 210 million per second.

10.3 Factorial representations of permutations

The factorial number system corresponds to the mixed radix bases $M = [2, 3, 4, \ldots]$ (*rising factorial basis*) or $M = [\ldots, 4, 3, 2]$ (*falling factorial basis*). A $(n-1)$-digit factorial number can have $n!$ different values. We develop different methods to convert factorial numbers to permutations and vice versa.

10.3.1 The Lehmer code

Each permutation of n distinct elements can be converted to a unique $(n-1)$-digit factorial number $A = [a_0, a_1, \ldots, a_{n-2}]$ in the falling factorial base by counting, for each index k, the number of elements with indices $j > k$ that are bigger than the current element [FXT: comb/fact2perm.cc]:

```
1   void perm2ffact(const ulong *x, ulong n, ulong *fc)
2   // Convert permutation in x[0,...,n-1] into
3   //   the (n-1) digit factorial representation in fc[0,...,n-2].
4   // One has: fc[0]<n, fc[1]<n-1, ... , fc[n-2]<2 (falling radices)
5   {
6       for (ulong k=0; k<n-1; ++k)
7       {
8           ulong xk = x[k];
9           ulong i = 0;
10          for (ulong j=k; j<n; ++j)  if ( x[j]<xk )  ++i;
11          fc[k] = i;
12      }
13  }
```

The routine works as long as all elements of the permutation are distinct. The factorial representation obtained by this method is called the *Lehmer code* of the permutation. For example, the permutation $[3, 0, 1, 4, 2]$ has the Lehmer code $[3, 0, 0, 1]$, because three elements smaller than the first element (3) lie right to it, no elements smaller than the second element (0) lies right to it, etc.

A routine that computes the permutation for a given Lehmer code is

```
1   void ffact2perm(const ulong *fc, ulong n, ulong *x)
2   // Inverse of perm2ffact():
3   // Convert the (n-1) digit factorial representation in fc[0,...,n-2].
4   // into permutation in x[0,...,n-1]
5   // Must have: fc[0]<n, fc[1]<n-1, ... , fc[n-2]<2 (falling radices)
6   {
7       for (ulong k=0; k<n; ++k)  x[k] = k;
8       for (ulong k=0; k<n-1; ++k)
9       {
10          ulong fa = fc[k];
11          if ( fa ) rotate_right1(x+k, fa+1);
12      }
13  }
```

A routine to compute the inverse permutation is

```
1   void ffact2invperm(const ulong *fc, ulong n, ulong *x)
2   // Convert the (n-1) digit factorial representation in fc[0,...,n-2].
3   // into permutation in x[0,...,n-1] such that
4   // the permutation is the inverse of the one computed via ffact2perm().
5   {
6       for (ulong k=0; k<n; ++k)  x[k] = n-1;  // "empty"
7       for (ulong k=0; k<n-1; ++k)
8       {
9           ulong fa = fc[k];
10          for (ulong j=0;  ; ++j)
11          {
12              if ( x[j]==n-1 )  // if empty
13              {
14                  if ( 0==fa )  { x[j] = k;  break; }
15                  --fa;
16              }
17          }
18      }
19  }
```

A similar method can compute a representation in the rising factorial base:

```
1   void perm2rfact(const ulong *x, ulong n, ulong *fc)
2   // Convert permutation in x[0,...,n-1] into
3   //   the (n-1) digit factorial representation in fc[0,...,n-2].
4   // One has: fc[0]<2, fc[1]<3, ... , fc[n-2]<n (rising radices)
5   {
6       for (ulong k=1; k<n; ++k)
7       {
```

	ffact	permutation	rev.compl.perm.	rfact
0:	[. . .]	[. 1 2 3]	[. 1 2 3]	[. . .]
1:	[1 . .]	[1 . 2 3]	[. 1 3 2]	[. . 1]
2:	[2 . .]	[2 . 1 3]	[. 2 3 1]	[. . 2]
3:	[3 . .]	[3 . 1 2]	[1 2 3 .]	[. . 3]
4:	[. 1 .]	[. 2 1 3]	[. 2 1 3]	[. 1 .]
5:	[1 1 .]	[1 2 . 3]	[. 3 1 2]	[. 1 1]
6:	[2 1 .]	[2 1 . 3]	[. 3 2 1]	[. 1 2]
7:	[3 1 .]	[3 1 . 2]	[1 3 2 .]	[. 1 3]
8:	[. 2 .]	[. 3 1 2]	[1 2 . 3]	[. 2 .]
9:	[1 2 .]	[1 3 . 2]	[1 3 . 2]	[. 2 1]
10:	[2 2 .]	[2 3 . 1]	[2 3 . 1]	[. 2 2]
11:	[3 2 .]	[3 2 . 1]	[2 3 1 .]	[. 2 3]
12:	[. . 1]	[. 1 3 2]	[1 . 2 3]	[1 . .]
13:	[1 . 1]	[1 . 3 2]	[1 . 3 2]	[1 . 1]
14:	[2 . 1]	[2 . 3 1]	[2 . 3 1]	[1 . 2]
15:	[3 . 1]	[3 . 2 1]	[2 1 3 .]	[1 . 3]
16:	[. 1 1]	[. 2 3 1]	[2 . 1 3]	[1 1 .]
17:	[1 1 1]	[1 2 3 .]	[3 . 1 2]	[1 1 1]
18:	[2 1 1]	[2 1 3 .]	[3 . 2 1]	[1 1 2]
19:	[3 1 1]	[3 1 2 .]	[3 1 2 .]	[1 1 3]
20:	[. 2 1]	[. 3 2 1]	[2 1 . 3]	[1 2 .]
21:	[1 2 1]	[1 3 2 .]	[3 1 . 2]	[1 2 1]
22:	[2 2 1]	[2 3 1 .]	[3 2 . 1]	[1 2 2]
23:	[3 2 1]	[3 2 1 .]	[3 2 1 .]	[1 2 3]

Figure 10.3-A: Numbers in falling factorial basis and permutations so that the number is the Lehmer code of it (left columns). Dots denote zeros. The rising factorial representation of the reversed and complemented permutation equals the reversed Lehmer code (right columns).

	rfact	permutation	rev.compl.perm.	ffact
0:	[. . .]	[. 1 2 3]	[. 1 2 3]	[. . .]
1:	[1 . .]	[1 . 2 3]	[. 1 3 2]	[. . 1]
2:	[. 1 .]	[. 2 1 3]	[. 2 1 3]	[. 1 .]
3:	[1 1 .]	[2 . 1 3]	[. 2 3 1]	[. 1 1]
4:	[. 2 .]	[1 2 . 3]	[. 3 1 2]	[. 2 .]
5:	[1 2 .]	[2 1 . 3]	[. 3 2 1]	[. 2 1]
6:	[. . 1]	[. 1 3 2]	[1 . 2 3]	[1 . .]
7:	[1 . 1]	[1 . 3 2]	[1 . 3 2]	[1 . 1]
8:	[. 1 1]	[. 3 1 2]	[1 2 . 3]	[1 1 .]
9:	[1 1 1]	[3 . 1 2]	[1 2 3 .]	[1 1 1]
10:	[. 2 1]	[1 3 . 2]	[1 3 . 2]	[1 2 .]
11:	[1 2 1]	[3 1 . 2]	[1 3 2 .]	[1 2 1]
12:	[. . 2]	[. 2 3 1]	[2 . 1 3]	[2 . .]
13:	[1 . 2]	[2 . 3 1]	[2 . 3 1]	[2 . 1]
14:	[. 1 2]	[. 3 2 1]	[2 1 . 3]	[2 1 .]
15:	[1 1 2]	[3 . 2 1]	[2 1 3 .]	[2 1 1]
16:	[. 2 2]	[2 3 . 1]	[2 3 . 1]	[2 2 .]
17:	[1 2 2]	[3 2 . 1]	[2 3 1 .]	[2 2 1]
18:	[. . 3]	[1 2 3 .]	[3 . 1 2]	[3 . .]
19:	[1 . 3]	[2 1 3 .]	[3 . 2 1]	[3 . 1]
20:	[. 1 3]	[1 3 2 .]	[3 1 . 2]	[3 1 .]
21:	[1 1 3]	[3 1 2 .]	[3 1 2 .]	[3 1 1]
22:	[. 2 3]	[2 3 1 .]	[3 2 . 1]	[3 2 .]
23:	[1 2 3]	[3 2 1 .]	[3 2 1 .]	[3 2 1]

Figure 10.3-B: Numbers in rising factorial basis and permutations so that the number is the Lehmer code of it (left columns). The reversed and complemented permutations and their falling factorial representations are shown in the right columns. They appear in lexicographic order.

```
8           ulong xk = x[k];
9           ulong i = 0;
10          for (ulong j=0; j<k; ++j)  if ( x[j]>xk )  ++i;
11          fc[k-1] = i;
12      }
13  }
```

Here we count, starting with the second element of the permutation, the number of elements to the left that are bigger than the current element. The inverse routine is

```
1   void rfact2perm(const ulong *fc, ulong n, ulong *x)
2   {
3       for (ulong k=0; k<n; ++k)  x[k] = k;
4       ulong *y = x+n;
5       for (ulong k=n-1;  k!=0;  --k, --y)
6       {
7           ulong fa = fc[k-1];
8           if ( fa )
9           {
10              ++fa;
11              rotate_left1(y-fa, fa);
12          }
13      }
14  }
```

A routine for the inverse permutation is

```
1   void rfact2invperm(const ulong *fc, ulong n, ulong *x)
2   // Convert the (n-1) digit factorial representation in fc[0,...,n-2].
3   //  into permutation in x[0,...,n-1] such that
4   // the permutation is the inverse of the one computed via rfact2perm().
5   {
6       for (ulong k=0; k<n; ++k)  x[k] = 0;  // "empty"
7       for (ulong k=n-2; (long)k>=0; --k)
8       {
9           ulong fa = fc[k];
10          for (ulong j=0;  ; ++j)
11          {
12              if ( x[j]==0 )  // if empty
13              {
14                  if ( 0==fa )  { x[j] = k+1;  break; }
15                  --fa;
16              }
17          }
18      }
19  }
```

The permutations corresponding to the Lehmer codes (in counting order) are shown in figure 10.3-A (left columns). The permutation whose rising factorial representation is the digit-reversed Lehmer code is obtained be reversing and complementing (replacing each element x by $n-1-x$) the original permutation:

```
Lehmer code  permutation     rev.perm    compl.rev.perm   rising fact
[3,0,0,1]    [3,0,1,4,2]    [2,4,1,0,3]  [2,0,3,4,1]      [1,0,0,3]
```

The permutations obtained from counting in the rising factorial base are shown in figure 10.3-B.

10.3.2 An representation via reversals

Replacing the rotations in the computation of a permutation from its Lehmer code by reversals one obtains a different one to one relation between factorial numbers and permutations. For the falling factorial basis one gets [FXT: comb/fact2perm-rev.cc]:

```
1   void perm2ffact_rev(const ulong *x, ulong n, ulong *fc)
2   {
3       ALLOCA(ulong, ti, n);  // inverse permutation
4       for (ulong k=0; k<n; ++k)  ti[x[k]] = k;
5       for (ulong k=0; k<n-1; ++k)
6       {
7           ulong j; // find element k
8           for (j=k; j<n; ++j)  if ( ti[j]==k )  break;
9           j -= k;
```

	ffact	permutation	inv.perm.	ffact
0:	[. . .]	[. 1 2 3]	[. 1 2 3]	[. . .]
1:	[1 . .]	[1 . 2 3]	[1 . 2 3]	[1 . .]
2:	[2 . .]	[2 1 . 3]	[2 1 . 3]	[2 . .]
3:	[3 . .]	[3 2 1 .]	[3 2 1 .]	[3 . .]
4:	[. 1 .]	[. 2 1 3]	[. 2 1 3]	[. 1 .]
5:	[1 1 .]	[1 2 . 3]	[2 . 1 3]	[2 1 .]
6:	[2 1 .]	[2 . 1 3]	[1 2 . 3]	[1 1 .]
7:	[3 1 .]	[3 1 2 .]	[3 1 2 .]	[3 1 .]
8:	[. 2 .]	[. 3 2 1]	[. 3 2 1]	[. 2 .]
9:	[1 2 .]	[1 3 2 .]	[3 . 2 1]	[3 2 1]
10:	[2 2 .]	[2 3 . 1]	[2 3 . 1]	[2 2 .]
11:	[3 2 .]	[3 . 1 2]	[1 2 3 .]	[1 1 1]
12:	[. . 1]	[. 1 3 2]	[. 1 3 2]	[. . 1]
13:	[1 . 1]	[1 . 3 2]	[1 . 3 2]	[1 . 1]
14:	[2 . 1]	[2 1 3 .]	[3 1 . 2]	[3 1 1]
15:	[3 . 1]	[3 2 . 1]	[2 3 1 .]	[2 2 1]
16:	[. 1 1]	[. 2 3 1]	[. 3 1 2]	[. 2 1]
17:	[1 1 1]	[1 2 3 .]	[3 . 1 2]	[3 2 .]
18:	[2 1 1]	[2 . 3 1]	[1 3 . 2]	[1 2 1]
19:	[3 1 1]	[3 1 . 2]	[2 1 3 .]	[2 . 1]
20:	[. 2 1]	[. 3 1 2]	[. 2 3 1]	[. 1 1]
21:	[1 2 1]	[1 3 . 2]	[2 . 3 1]	[2 1 1]
22:	[2 2 1]	[2 3 1 .]	[3 2 . 1]	[3 . 1]
23:	[3 2 1]	[3 . 2 1]	[1 3 2 .]	[1 2 .]

Figure 10.3-C: Numbers in falling factorial basis and permutations so that the number is the alternative (reversal-) code of it (left columns). The inverse permutations and their rising factorial representations are shown in the right columns.

	rfact	permutation	inv.perm.	rfact
0:	[. . .]	[. 1 2 3]	[. 1 2 3]	[. . .]
1:	[1 . .]	[1 . 2 3]	[1 . 2 3]	[1 . .]
2:	[. 1 .]	[. 2 1 3]	[. 2 1 3]	[. 1 .]
3:	[1 1 .]	[2 . 1 3]	[1 2 . 3]	[1 2 .]
4:	[. 2 .]	[2 1 . 3]	[2 1 . 3]	[. 2 .]
5:	[1 2 .]	[1 2 . 3]	[2 . 1 3]	[1 1 .]
6:	[. . 1]	[. 1 3 2]	[. 1 3 2]	[. . 1]
7:	[1 . 1]	[1 . 3 2]	[1 . 3 2]	[1 . 1]
8:	[. 1 1]	[. 3 1 2]	[. 2 3 1]	[. 1 2]
9:	[1 1 1]	[3 . 1 2]	[1 2 3 .]	[. 2 3]
10:	[. 2 1]	[3 1 . 2]	[2 1 3 .]	[1 2 3]
11:	[1 2 1]	[1 3 . 2]	[2 . 3 1]	[1 1 2]
12:	[. . 2]	[. 3 2 1]	[. 3 2 1]	[. . 2]
13:	[1 . 2]	[3 . 2 1]	[1 3 2 .]	[1 1 3]
14:	[. 1 2]	[. 2 3 1]	[. 3 1 2]	[. 1 1]
15:	[1 1 2]	[2 . 3 1]	[1 3 . 2]	[1 2 1]
16:	[. 2 2]	[2 3 . 1]	[2 3 . 1]	[. 2 2]
17:	[1 2 2]	[3 2 . 1]	[2 3 1 .]	[1 . 3]
18:	[. . 3]	[3 2 1 .]	[3 2 1 .]	[. . 3]
19:	[1 . 3]	[2 3 1 .]	[3 2 . 1]	[1 2 2]
20:	[. 1 3]	[3 1 2 .]	[3 1 2 .]	[. 1 3]
21:	[1 1 3]	[1 3 2 .]	[3 . 2 1]	[1 . 2]
22:	[. 2 3]	[1 2 3 .]	[3 . 1 2]	[1 1 1]
23:	[1 2 3]	[2 1 3 .]	[3 1 . 2]	[. 2 1]

Figure 10.3-D: Numbers in rising factorial basis and permutations so that the number is the alternative (reversal-) code of it (left columns). The inverse permutations and their falling factorial representations are shown in the right columns.

```
10          fc[k] = j;
11          reverse(ti+k, j+1);
12      }
13  }
```

The routine is the inverse of

```
1   void ffact2perm_rev(const ulong *fc, ulong n, ulong *x)
2   {
3       for (ulong k=0; k<n; ++k)  x[k] = k;
4       for (ulong k=0; k<n-1; ++k)
5       {
6           ulong fa = fc[k];
7           // Lehmer: rotate_right1(x+k, fa+1);
8           if ( fa ) reverse(x+k, fa+1);
9       }
10  }
```

Figure 10.3-C shows the permutations of 4 elements and their falling factorial representations, it was created with the program [FXT: comb/fact2perm-rev-demo.cc]. The routines for the rising factorial (figure 10.3-D) basis are

```
1   void perm2rfact_rev(const ulong *x, ulong n, ulong *fc)
2   {
3       ALLOCA(ulong, ti, n);  // inverse permutation
4       for (ulong k=0; k<n; ++k)  ti[x[k]] = k;
5       for (ulong k=n-1; k!=0; --k)
6       {
7           ulong j; // find element k
8           for (j=0; j<=k; ++j)  if ( ti[j]==k )  break;
9           j = k - j;
10          fc[k-1] = j;
11          reverse(ti+k-j, j+1);
12      }
13  }
```

and

```
1   void rfact2perm_rev(const ulong *fc, ulong n, ulong *x)
2   {
3       for (ulong k=0; k<n; ++k)  x[k] = k;
4       ulong *y = x+n;
5       for (ulong k=n-1;  k!=0;  --k, --y)
6       {
7           ulong fa = fc[k-1];
8           if ( fa )
9           {
10              ++fa;
11              // Lehmer: rotate_left1(y-fa, fa);
12              reverse(y-fa, fa);
13          }
14      }
15  }
```

10.3.3 A representation via swaps

The routines for the conversion from permutations to factorial representations shown so far have complexity n^2. The following routines compute different factorial representations with complexity n [FXT: comb/fact2perm-swp.cc]:

```
1   void
2   perm2ffact_swp(const ulong *x, ulong n, ulong *fc)
3   // Convert permutation in x[0,...,n-1] into
4   //   the (n-1) digit (swaps-) factorial representation in fc[0,...,n-2].
5   // One has: fc[0]<n, fc[1]<n-1, ... , fc[n-2]<2 (falling radices)
6   {
7       ALLOCA(ulong, t, n);
8       for (ulong k=0; k<n; ++k)  t[k] = x[k];
9       ALLOCA(ulong, ti, n);  // inverse permutation
10      for (ulong k=0; k<n; ++k)  ti[t[k]] = k;
11
12      for (ulong k=0; k<n-1; ++k)
```

ffact.	permutation	inv.perm.	rfact.
[. . .]	[. 1 2 3]	[. 1 2 3]	[. . .]
[1 . .]	[1 . 2 3]	[1 . 2 3]	[. . 1]
[2 . .]	[2 1 . 3]	[2 1 . 3]	[. . 2]
[3 . .]	[3 1 2 .]	[3 1 2 .]	[. . 3]
[. 1 .]	[. 2 1 3]	[. 2 1 3]	[. 1 .]
[1 1 .]	[1 2 . 3]	[2 . 1 3]	[. 1 1]
[2 1 .]	[2 . 1 3]	[1 2 . 3]	[. 1 2]
[3 1 .]	[3 2 1 .]	[3 2 1 .]	[. 1 3]
[. 2 .]	[. 3 2 1]	[. 3 2 1]	[. 2 .]
[1 2 .]	[1 3 2 .]	[3 . 2 1]	[. 2 1]
[2 2 .]	[2 3 . 1]	[2 3 . 1]	[. 2 2]
[3 2 .]	[3 . 2 1]	[1 3 2 .]	[. 2 3]
[. . 1]	[. 1 3 2]	[. 1 3 2]	[1 . .]
[1 . 1]	[1 . 3 2]	[1 . 3 2]	[1 . 1]
[2 . 1]	[2 1 3 .]	[3 1 . 2]	[1 . 2]
[3 . 1]	[3 1 . 2]	[2 1 3 .]	[1 . 3]
[. 1 1]	[. 2 3 1]	[. 3 1 2]	[1 1 .]
[1 1 1]	[1 2 3 .]	[3 . 1 2]	[1 1 1]
[2 1 1]	[2 . 3 1]	[1 3 . 2]	[1 1 2]
[3 1 1]	[3 2 . 1]	[2 3 1 .]	[1 1 3]
[. 2 1]	[. 3 1 2]	[. 2 3 1]	[1 2 .]
[1 2 1]	[1 3 . 2]	[2 . 3 1]	[1 2 1]
[2 2 1]	[2 3 1 .]	[3 2 . 1]	[1 2 2]
[3 2 1]	[3 . 1 2]	[1 2 3 .]	[1 2 3]

Figure 10.3-E: Numbers in falling factorial basis and permutations so that the number is the alternative (swaps-) code of it (left columns). The inverse permutations and their rising factorial representations are shown in the right columns.

rfact	permutation	inv.perm.	ffact
[. . .]	[. 1 2 3]	[. 1 2 3]	[. . .]
[1 . .]	[. 1 3 2]	[. 1 3 2]	[. . 1]
[. 1 .]	[. 2 1 3]	[. 2 1 3]	[. 1 .]
[1 1 .]	[. 3 1 2]	[. 2 3 1]	[. 1 1]
[. 2 .]	[. 3 2 1]	[. 3 2 1]	[. 2 .]
[1 2 .]	[. 2 3 1]	[. 3 1 2]	[. 2 1]
[. . 1]	[1 . 2 3]	[1 . 2 3]	[1 . .]
[1 . 1]	[1 . 3 2]	[1 . 3 2]	[1 . 1]
[. 1 1]	[2 . 1 3]	[1 2 . 3]	[1 1 .]
[1 1 1]	[3 . 1 2]	[1 2 3 .]	[1 1 1]
[. 2 1]	[3 . 2 1]	[1 3 2 .]	[1 2 .]
[1 2 1]	[2 . 3 1]	[1 3 . 2]	[1 2 1]
[. . 2]	[2 1 . 3]	[2 1 . 3]	[2 . .]
[1 . 2]	[3 1 . 2]	[2 1 3 .]	[2 . 1]
[. 1 2]	[1 2 . 3]	[2 . 1 3]	[2 1 .]
[1 1 2]	[1 3 . 2]	[2 . 3 1]	[2 1 1]
[. 2 2]	[2 3 . 1]	[2 3 . 1]	[2 2 .]
[1 2 2]	[3 2 . 1]	[2 3 1 .]	[2 2 1]
[. . 3]	[3 1 2 .]	[3 1 2 .]	[3 . .]
[1 . 3]	[2 1 3 .]	[3 1 . 2]	[3 . 1]
[. 1 3]	[3 2 1 .]	[3 2 1 .]	[3 1 .]
[1 1 3]	[2 3 1 .]	[3 2 . 1]	[3 1 1]
[. 2 3]	[1 3 2 .]	[3 . 2 1]	[3 2 .]
[1 2 3]	[1 2 3 .]	[3 . 1 2]	[3 2 1]

Figure 10.3-F: Numbers in rising factorial basis and permutations so that the number is the alternative (swaps-) code of it (left columns). The inverse permutations and their falling factorial representations are shown in the right columns.

```
13      {
14          ulong tk = t[k];  // >= k
15          fc[k] = tk - k;
16          ulong j = ti[k];  // location of element k
17          ti[tk] = ti[t[j]];
18          t[j] = tk;
19      }
20  }
```

```
1   void perm2rfact_swp(const ulong *x, ulong n, ulong *fc)
2   // Convert permutation in x[0,...,n-1] into
3   //   the (n-1) digit (swaps-) factorial representation in fc[0,...,n-2].
4   // One has: fc[0]<2, fc[1]<3, ... , fc[n-2]<n (rising radices)
5   {
6       ALLOCA(ulong, t, n);
7       for (ulong k=0; k<n; ++k)  t[k] = x[k];
8       ALLOCA(ulong, ti, n);  // inverse permutation
9       for (ulong k=0; k<n; ++k)  ti[t[k]] = k;
10
11      for (ulong k=0; k<n-1; ++k)
12      {
13          ulong j = ti[k];  // location of element k, j>=k
14          fc[n-2-k] = j - k;
15          ulong tk = t[k];
16          ti[tk] = ti[t[j]];
17          t[j] = tk;
18      }
19  }
```

Their inverses also have complexity n. The routine for falling base is

```
1   void ffact2perm_swp(const ulong *fc, ulong n, ulong *x)
2   // Inverse of perm2ffact_swp().
3   {
4       for (ulong k=0; k<n; ++k)  x[k] = k;
5       for (ulong k=0; k<n-1; ++k)
6       {
7           ulong fa = fc[k];
8           swap2( x[k], x[k+fa] );
9       }
10  }
```

The routine for the rising base is

```
1   void rfact2perm_swp(const ulong *fc, ulong n, ulong *x)
2   // Inverse of perm2rfact_swp().
3   {
4       for (ulong k=0; k<n; ++k)  x[k] = k;
5       for (ulong k=0,j=n-2; k<n-1; ++k,--j)
6       {
7           ulong fa = fc[k];
8           swap2( x[j], x[j+fa] );
9       }
10  }
```

The permutations corresponding to the alternative codes for the falling basis are shown in figure 10.3-E (left columns). The inverse permutation has the rising factorial representation that is digit-reversed representation (right columns). The permutations corresponding to the alternative codes for rising basis are shown in figure 10.3-F The listings were created with the program [FXT: comb/fact2perm-swp-demo.cc]. The routines can serve as a means to find interesting orders of permutations. Indeed, the permutation generator shown in section 10.4 on page 245 was found this way. A recursive algorithm for the (inverse) permutations shown in figure 10.3-F is given in section 10.13 on page 278.

10.3.4 Cyclic permutations

Cyclic permutations of n elements (permutations that consists of one cycle of size n, see section 2.11.4 on page 116) can be obtained from length-$(n-2)$ factorial numbers. We give routines for both falling and rising base [FXT: comb/fact2cyclic.cc]:

falling fact.	permutation	cycle	inv.perm.
[. . .]	[1 2 3 4 0]	(0, 1, 2, 3, 4,)	[4 0 1 2 3]
[1 . .]	[4 2 3 0 1]	(0, 4, 1, 2, 3,)	[3 4 1 2 0]
[2 . .]	[1 4 3 0 2]	(0, 1, 4, 2, 3,)	[3 0 4 2 1]
[3 . .]	[1 2 4 0 3]	(0, 1, 2, 4, 3,)	[3 0 1 4 2]
[. 1 .]	[3 2 4 1 0]	(0, 3, 1, 2, 4,)	[4 3 1 0 2]
[1 1 .]	[3 2 0 4 1]	(0, 3, 4, 1, 2,)	[2 4 1 0 3]
[2 1 .]	[3 4 0 1 2]	(0, 3, 1, 4, 2,)	[2 3 4 0 1]
[3 1 .]	[4 2 0 1 3]	(0, 4, 3, 1, 2,)	[2 3 1 4 0]
[. 2 .]	[1 3 4 2 0]	(0, 1, 3, 2, 4,)	[4 0 3 1 2]
[1 2 .]	[4 3 0 2 1]	(0, 4, 1, 3, 2,)	[2 4 3 1 0]
[2 2 .]	[1 3 0 4 2]	(0, 1, 3, 4, 2,)	[2 0 4 1 3]
[3 2 .]	[1 4 0 2 3]	(0, 1, 4, 3, 2,)	[2 0 3 4 1]
[. . 1]	[2 3 1 4 0]	(0, 2, 1, 3, 4,)	[4 2 0 1 3]
[1 . 1]	[2 3 4 0 1]	(0, 2, 4, 1, 3,)	[3 4 0 1 2]
[2 . 1]	[4 3 1 0 2]	(0, 4, 2, 1, 3,)	[3 2 4 1 0]
[3 . 1]	[2 4 1 0 3]	(0, 2, 1, 4, 3,)	[3 2 0 4 1]
[. 1 1]	[2 4 3 1 0]	(0, 2, 3, 1, 4,)	[4 3 0 2 1]
[1 1 1]	[2 0 3 4 1]	(0, 2, 3, 4, 1,)	[1 4 0 2 3]
[2 1 1]	[4 0 3 1 2]	(0, 4, 2, 3, 1,)	[1 3 4 2 0]
[3 1 1]	[2 0 4 1 3]	(0, 2, 4, 3, 1,)	[1 3 0 4 2]
[. 2 1]	[3 4 1 2 0]	(0, 3, 2, 1, 4,)	[4 2 3 0 1]
[1 2 1]	[3 0 4 2 1]	(0, 3, 2, 4, 1,)	[1 4 3 0 2]
[2 2 1]	[3 0 1 4 2]	(0, 3, 4, 2, 1,)	[1 2 4 0 3]
[3 2 1]	[4 0 1 2 3]	(0, 4, 3, 2, 1,)	[1 2 3 4 0]

Figure 10.3-G: Numbers in falling factorial basis and the corresponding cyclic permutations.

rising fact.	permutation	cycle	inv.perm.
[. . .]	[1 2 3 4 0]	(0, 1, 2, 3, 4,)	[4 0 1 2 3]
[1 . .]	[2 3 1 4 0]	(0, 2, 1, 3, 4,)	[4 2 0 1 3]
[. 1 .]	[3 2 4 1 0]	(0, 3, 1, 2, 4,)	[4 3 1 0 2]
[1 1 .]	[2 4 3 1 0]	(0, 2, 3, 1, 4,)	[4 3 0 2 1]
[. 2 .]	[1 3 4 2 0]	(0, 1, 3, 2, 4,)	[4 0 3 1 2]
[1 2 .]	[3 4 1 2 0]	(0, 3, 2, 1, 4,)	[4 2 3 0 1]
[. . 1]	[4 2 3 0 1]	(0, 4, 1, 2, 3,)	[3 4 1 2 0]
[1 . 1]	[2 3 4 0 1]	(0, 2, 4, 1, 3,)	[3 4 0 1 2]
[. 1 1]	[3 2 0 4 1]	(0, 3, 4, 1, 2,)	[2 4 1 0 3]
[1 1 1]	[2 0 3 4 1]	(0, 2, 3, 4, 1,)	[1 4 0 2 3]
[. 2 1]	[4 3 0 2 1]	(0, 4, 1, 3, 2,)	[2 4 3 1 0]
[1 2 1]	[3 0 4 2 1]	(0, 3, 2, 4, 1,)	[1 4 3 0 2]
[. . 2]	[1 4 3 0 2]	(0, 1, 4, 2, 3,)	[3 0 4 2 1]
[1 . 2]	[4 3 1 0 2]	(0, 4, 2, 1, 3,)	[3 2 4 1 0]
[. 1 2]	[3 4 0 1 2]	(0, 3, 1, 4, 2,)	[2 3 4 0 1]
[1 1 2]	[4 0 3 1 2]	(0, 4, 2, 3, 1,)	[1 3 4 2 0]
[. 2 2]	[1 3 0 4 2]	(0, 1, 3, 4, 2,)	[2 0 4 1 3]
[1 2 2]	[3 0 1 4 2]	(0, 3, 4, 2, 1,)	[1 2 4 0 3]
[. . 3]	[1 2 4 0 3]	(0, 1, 2, 4, 3,)	[3 0 1 4 2]
[1 . 3]	[2 4 1 0 3]	(0, 2, 1, 4, 3,)	[3 2 0 4 1]
[. 1 3]	[4 2 0 1 3]	(0, 4, 3, 1, 2,)	[2 3 1 4 0]
[1 1 3]	[2 0 4 1 3]	(0, 2, 4, 3, 1,)	[1 3 0 4 2]
[. 2 3]	[1 4 0 2 3]	(0, 1, 4, 3, 2,)	[2 0 3 4 1]
[1 2 3]	[4 0 1 2 3]	(0, 4, 3, 2, 1,)	[1 2 3 4 0]

Figure 10.3-H: Numbers in rising factorial basis and corresponding cyclic permutations.

```
1   void ffact2cyclic(const ulong *fc, ulong n, ulong *x)
2   // Generate cyclic permutation (standard representation) in x[]
3   //   from the (n-2) digit factorial number in fc[0,...,n-3].
4   // Falling radices:  [n-1, ..., 3, 2]
5   {
6       for (ulong k=0; k<n; ++k)  x[k] = k;
7
8       for (ulong k=n-1; k>1; --k)
9       {
10          ulong z = n-1-k; // 0, ..., n-3
11          ulong i = fc[z];
12          swap2(x[k], x[i]);
13      }
14
15      if ( n>1 )  swap2(x[0], x[1]);
16  }
```

```
1   void rfact2cyclic(const ulong *fc, ulong n, ulong *x)
2   // Generate cyclic permutation (standard representation) in x[]
3   //   from the (n-2) digit factorial number in fc[0,...,n-3].
4   // Rising radices:  [2, 3, ..., n-1]
5   {
6       for (ulong k=0; k<n; ++k)  x[k] = k;
7
8       for (ulong k=n-1; k>1; --k)
9       {
10          ulong i = fc[k-2];  // k-2 == n-3, ..., 0
11          swap2(x[k], x[i]);
12      }
13
14      if ( n>1 )  swap2(x[0], x[1]);
15  }
```

The cyclic permutations of 5 elements are shown in figures 10.3-G (falling base) and 10.3-H (rising base). The listings were created with the program [FXT: comb/fact2cyclic-demo.cc]. Note that the cycle representation could be obtained by applying the transformations to (all) permutations to all but the first element. That is, one can generate all cyclic permutations in cycle form by permuting all elements but the first with any permutation algorithm.

10.4 An order from reversing prefixes

A surprisingly simple algorithm for the generation of all permutations is obtained by mixed radix counting with the radices [2, 3, 4, ...] (column `digits` in figure 10.4-A). Whenever the first j digits change with an increment then the permutation is updated by reversing the first $j+1$ elements. As with lex order the first half of the permutations are the complements of the permutations in the second half, now rewrite relation 10.1-1 on page 233 as

$$R \;=\; I\,C\,I \tag{10.4-1}$$

to see that the first half of the inverse permutations are the reversed inverse permutations in the second half. This can (for n even) also be observed from the positions of the largest element in the inverse permutations. An implementation is [FXT: `class perm_rev` in comb/perm-rev.h]:

```
1   class perm_rev
2   {
3   public:
4       ulong *d_;  // mixed radix digits with radix = [2, 3, 4, ..., n-1, (sentinel=-1)]
5       ulong *p_;  // permutation
6       ulong n_;   // permutations of n elements
7
8   public:
9       perm_rev(ulong n)
10      {
11          n_ = n;
12          p_ = new ulong[n_];
13          d_ = new ulong[n_];
14          d_[n-1] = -1UL;  // sentinel
```

```
              permutation     rfact           inv. perm.
        0:    [ . 1 2 3 ]    [ . . . ]        [ . 1 2 3 ]
        1:    [ 1 . 2 3 ]    [ 1 . . ]        [ 1 . 2 3 ]
        2:    [ 2 . 1 3 ]    [ . 1 . ]        [ 1 2 . 3 ]
        3:    [ . 2 1 3 ]    [ 1 1 . ]        [ . 2 1 3 ]
        4:    [ 1 2 . 3 ]    [ . 2 . ]        [ 2 . 1 3 ]
        5:    [ 2 1 . 3 ]    [ 1 2 . ]        [ 2 1 . 3 ]
        6:    [ 3 . 1 2 ]    [ . . 1 ]        [ 1 2 3 . ]
        7:    [ . 3 1 2 ]    [ 1 . 1 ]        [ . 2 3 1 ]
        8:    [ 1 3 . 2 ]    [ . 1 1 ]        [ 2 . 3 1 ]
        9:    [ 3 1 . 2 ]    [ 1 1 1 ]        [ 2 1 3 . ]
       10:    [ . 1 3 2 ]    [ . 2 1 ]        [ . 1 3 2 ]
       11:    [ 1 . 3 2 ]    [ 1 2 1 ]        [ 1 . 3 2 ]
       12:    [ 2 3 . 1 ]    [ . . 2 ]        [ 2 3 . 1 ]
       13:    [ 3 2 . 1 ]    [ 1 . 2 ]        [ 2 3 1 . ]
       14:    [ . 2 3 1 ]    [ . 1 2 ]        [ . 3 1 2 ]
       15:    [ 2 . 3 1 ]    [ 1 1 2 ]        [ 1 3 . 2 ]
       16:    [ 3 . 2 1 ]    [ . 2 2 ]        [ 1 3 2 . ]
       17:    [ . 3 2 1 ]    [ 1 2 2 ]        [ . 3 2 1 ]
       18:    [ 1 2 3 . ]    [ . . 3 ]        [ 3 . 1 2 ]
       19:    [ 2 1 3 . ]    [ 1 . 3 ]        [ 3 1 . 2 ]
       20:    [ 3 1 2 . ]    [ . 1 3 ]        [ 3 1 2 . ]
       21:    [ 1 3 2 . ]    [ 1 1 3 ]        [ 3 . 2 1 ]
       22:    [ 2 3 1 . ]    [ . 2 3 ]        [ 3 2 . 1 ]
       23:    [ 3 2 1 . ]    [ 1 2 3 ]        [ 3 2 1 . ]
```

Figure 10.4-A: All permutations of 4 elements in an order where the first $j+1$ elements are reversed when the first j digits change in the mixed radix counting sequence with radices [2, 3, 4, ...].

```
15          first();
16      }
17
18      ~perm_rev()
19      {
20          delete [] p_;
21          delete [] d_;
22      }
23
24      void first()
25      {
26          for (ulong k=0; k<n_-1; ++k)  d_[k] = 0;
27          for (ulong k=0; k<n_; ++k)  p_[k] = k;
28      }
29
30      void last()
31      {
32          for (ulong k=0; k<n_-1; ++k)  d_[k] = k+1;
33          for (ulong k=0; k<n_; ++k)  p_[k] = n_-1-k;
34      }
```

The update routines are quite concise:

```
1       bool next()
2       {
3           // increment mixed radix number:
4           ulong j = 0;
5           while ( d_[j]==j+1 )  { d_[j]=0; ++j; }
6
7           // j==n-1 for last permutation
8           if ( j!=n_-1 ) // only if no overflow
9           {
10              ++d_[j];
11              reverse(p_, j+2); // update permutation
12              return true;
13          }
14          else  return false;
15      }
16
17      bool prev()
18      {
19          // decrement mixed radix number:
20          ulong j = 0;
21          while ( d_[j]==0 )  { d_[j]=j+1; ++j; }
```

```
22
23          // j==n-1 for last permutation
24          if ( j!=n_-1 ) // only if no overflow
25          {
26              --d_[j];
27              reverse(p_, j+2); // update permutation
28              return true;
29          }
30          else  return false;
31      }
32  };
```

Note that the routines work for arbitrary (distinct) entries of the array `p_[]`.

An upper bound for the average number of elements that are moved in the transitions when generating all $N = n!$ permutations is $e \approx 2.7182818$ so the algorithm is CAT. The implementation is actually fast, it generates more than 110 million permutations per second. Usage of the class is as simple as:

```
ulong n = 4;  // Number of elements to permute
perm_rev P(n);
P.first();
do
{
    // Use permutation here
}
while ( P.next() );
```

Figure 10.4-A was produced with [FXT: comb/perm-rev-demo.cc]. The described method is given in [279].

10.4.1 Optimizing the update routine

One can optimize the update routine by observing that 5 out of six updates are the swaps

```
(0,1)  (0,2)  (0,1)  (0,2)  (0,1)
```

We use a counter `ct_` and modify the methods `first()` and `next()` [FXT: `class perm_rev2` in comb/perm-rev2.h]:

```
1   class perm_rev2
2   {
3     [--snip--]
4       void first()
5       {
6           for (ulong k=0; k<n_-1; ++k)  d_[k] = 0;
7           for (ulong k=0; k<n_; ++k)  p_[k] = k;
8           ct_ = 5;
9       }
10
11      bool next()
12      {
13          if ( ct_!=0 )  // easy case(s)
14          {
15              --ct_;
16              swap2(p_[0], p_[1 + (ct_ & 1)]);
17              return true;
18          }
19          else  // increment mixed radix number:
20          {
21              ct_ = 5;  // reset counter
22
23              ulong j = 2;  // note: start with 2
24              while ( d_[j]==j+1 )  { d_[j]=0; ++j; }
25
26              // j==n-1 for last permutation
27              if ( j!=n_-1 ) // only if no overflow
28              {
29                  ++d_[j];
30                  reverse(p_, j+2); // update permutation
31                  return true;
32              }
33              else  return false;
34          }
35      }
36    [--snip--]
```

The speedup is remarkable, about 186 million permutations per second can be generated (about 11.8 cycles per update). If arrays are used instead of pointers, the rate is about 200 million per second (10.8 cycles per update). The routine can be used for permutations of at least three elements [FXT: comb/perm-rev2-demo.cc].

10.4.2 Method for unranking

Conversion of a mixed radix (rising factorial) number into the corresponding permutation proceeds as exemplified for the 16-th permutation ($15 = 1 \cdot 1 + 1 \cdot 2 + 2 \cdot 6$, so `d=[1,1,2]`):

```
 1:   p=[ 0, 1, 2, 3 ]   d=[ 0, 0, 0 ] // start
13:   p=[ 2, 3, 0, 1 ]   d=[ 0, 0, 2 ] // right rotate all elements twice
15:   p=[ 0, 2, 3, 1 ]   d=[ 0, 1, 2 ] // right rotate first three elements
16:   p=[ 2, 0, 3, 1 ]   d=[ 1, 1, 2 ] // right rotate first two elements
```

The idea can be implemented as

```
1       void goto_rfact(const ulong *d)
2       // Goto permutation corresponding to d[] (i.e. unrank d[]).
3       // d[] must be a valid (rising) factorial mixed radix string:
4       // d[]==[d(0), d(1), d(2), ..., d(n-2)]  (n-1 elements) where 0<=d(j)<=j+1
5       {
6           for (ulong k=0; k<n_; ++k)  p_[k] = k;
7           for (ulong k=0; k<n_-1; ++k)  d_[k] = d[k];
8           for (long j=n_-2; j>=0; --j)  rotate_right(p_, j+2, d_[j]);
9       }
```

10.5 Minimal-change order (Heap's algorithm)

	permutation	swap	digits	rfact(perm)	inv. perm.
0:	[. 1 2 3]	(0, 0)	[. . .]	[. . .]	[. 1 2 3]
1:	[1 . 2 3]	(1, 0)	[1 . .]	[1 . .]	[1 . 2 3]
2:	[2 . 1 3]	(2, 0)	[. 1 .]	[1 1 .]	[1 2 . 3]
3:	[. 2 1 3]	(1, 0)	[1 1 .]	[. 1 .]	[. 2 1 3]
4:	[1 2 . 3]	(2, 0)	[. 2 .]	[. 2 .]	[2 . 1 3]
5:	[2 1 . 3]	(1, 0)	[1 2 .]	[1 2 .]	[2 1 . 3]
6:	[3 1 . 2]	(3, 0)	[. . 1]	[1 2 1]	[2 1 3 .]
7:	[1 3 . 2]	(1, 0)	[1 . 1]	[. 2 1]	[2 . 3 1]
8:	[. 3 1 2]	(2, 0)	[. 1 1]	[. 1 1]	[. 2 3 1]
9:	[3 . 1 2]	(1, 0)	[1 1 1]	[1 1 1]	[1 2 3 .]
10:	[1 . 3 2]	(2, 0)	[. 2 1]	[1 . 1]	[1 . 3 2]
11:	[. 1 3 2]	(1, 0)	[1 2 1]	[. . 1]	[. 1 3 2]
12:	[. 2 3 1]	(3, 1)	[. . 2]	[. . 2]	[. 3 1 2]
13:	[2 . 3 1]	(1, 0)	[1 . 2]	[1 . 2]	[1 3 . 2]
14:	[3 . 2 1]	(2, 0)	[. 1 2]	[1 1 2]	[1 3 2 .]
15:	[. 3 2 1]	(1, 0)	[1 1 2]	[. 1 2]	[. 3 2 1]
16:	[2 3 . 1]	(2, 0)	[. 2 2]	[. 2 2]	[2 3 . 1]
17:	[3 2 . 1]	(1, 0)	[1 2 2]	[1 2 2]	[2 3 1 .]
18:	[3 2 1 .]	(3, 2)	[. . 3]	[1 2 3]	[3 2 1 .]
19:	[2 3 1 .]	(1, 0)	[1 . 3]	[. 2 3]	[3 2 . 1]
20:	[1 3 2 .]	(2, 0)	[. 1 3]	[. 1 3]	[3 . 2 1]
21:	[3 1 2 .]	(1, 0)	[1 1 3]	[1 1 3]	[3 1 2 .]
22:	[2 1 3 .]	(2, 0)	[. 2 3]	[1 . 3]	[3 1 . 2]
23:	[1 2 3 .]	(1, 0)	[1 2 3]	[. . 3]	[3 . 1 2]

Figure 10.5-A: The permutations of 4 elements in a minimal-change order. Dots denote zeros.

Figure 10.5-A shows the permutations of 4 elements in a *minimal-change order*: just two elements are swapped with each update. The column labeled `digits` shows the mixed radix numbers (rising factorial base, see section 10.3 on page 236) in counting order. Let j be the position of the rightmost change of the mixed radix string R. Then the swap is $(j+1, x)$ where $x = 0$ if j is odd, and $x = R_j - 1$ if j is even. The sequence of values $j + 1$ starts

```
1, 2, 1, 2, 1, 3, 1, 2, 1, 2, 1, 3, 1, 2, 1, 2, 1, 3, 1, 2, 1, 2, 1, 4, 1, 2, 1, ...
```

The n-th value (starting with $n = 1$) is the largest z such that $z!$ divides n (entry A055881 of [240]). The column labeled `rfact(perm)` of the figure shows the rising factorial representations of the permutation, see section 10.8 on page 258. The column is a Gray code only for permutations of up to four elements.

An implementation of the algorithm (given in [143]) is [FXT: `class perm_heap` in comb/perm-heap.h]:

```
class perm_heap
{
public:
    ulong *d_;  // mixed radix digits with radix = [2, 3, 4, ..., n-1, (sentinel=-1)]
    ulong *p_;  // permutation
    ulong n_;   // permutations of n elements
    ulong sw1_, sw2_;   // indices of swapped elements
  [--snip--]
```

The computation of the successor is simple:

```
    bool next()
    {
        // increment mixed radix number:
        ulong j = 0;
        while ( d_[j]==j+1 )  { d_[j]=0; ++j; }

        // j==n-1 for last permutation
        if ( j!=n_-1 ) // only if no overflow
        {
            ulong k = j+1;
            ulong x = ( (k&1)  ? d_[j] : 0);
            swap2(p_[k], p_[x]);
            sw1_ = k;  sw2_ = x;

            ++d_[j];
            return true;
        }
        else  return false;
    }
  [--snip--]
```

About 115 million permutations are generated per second. Often one will only use the indices of the swapped elements to update the visited configurations:

```
void get_swap(ulong &s1, ulong &s2)  const  { s1=sw1_; s2=sw2_; }
```

Then the statement `swap2(p_[k], p_[x]);` in the update routine can be omitted which leads to a rate of 165 million permutations per second. Figure 10.5-A shows the permutations of 4 elements, it was created with the program [FXT: comb/perm-heap-demo.cc].

10.5.1 Optimized implementation

The algorithm can be optimized by treating 5 out of 6 cases separately, those where the first or second digit in the mixed radix number changes. We use a counter `ct_` that is decremented [FXT: `class perm_heap2` in comb/perm-heap2.h]:

```
class perm_heap2
{
public:
    ulong *d_;  // mixed radix digits with radix = [2, 3, 4, 5, ..., n-1, (sentinel=-1)]
    ulong *p_;  // permutation
    ulong n_;   // permutations of n elements
    ulong sw1_, sw2_;   // indices of swapped elements
    ulong ct_;  // count 5,4,3,2,1,(0); nonzero ==> easy cases
  [--snip--]
```

The counter is set to 5 in the method `first()`. The update routine is

```
    bool next()
    {
        if ( ct_!=0 )  // easy case(s)
        {
```

```
5               --ct_;
6               sw1_ = 1 + (ct_ & 1);  // == 1,2,1,2,1
7               sw2_ = 0;
8               swap2(p_[sw1_], p_[sw2_]);
9               return true;
10          }
11          else
12          {
13              ct_ = 5;  // reset counter
14
15              // increment mixed radix number:
16              ulong j = 2;
17              while ( d_[j]==j+1 )  { d_[j]=0; ++j; }
18
19              // j==n-1 for last permutation
20              if ( j!=n_-1 ) // only if no overflow
21              {
22                  ulong k = j+1;
23                  ulong x = ( (k&1)  ? d_[j] : 0);
24                  sw1_ = k;  sw2_ = x;
25                  swap2(p_[sw1_], p_[sw2_]);
26                  ++d_[j];
27                  return true;
28              }
29              else  return false;
30          }
31      }
```

Note that the routine only works for permutations of at least three elements. Usage of the class is shown in [FXT: comb/perm-heap2-demo.cc]. The rate is about 190 million updates per second.

10.5.2 Computing just the swaps

If only the swaps are of interest we can simply omit all statements involving the permutation array `p_[]`. The implementation is [FXT: `class perm_heap2_swaps` in comb/perm-heap2-swaps.h], usage of the class is shown in [FXT: comb/perm-heap2-swaps-demo.cc]:

```
1       ulong n = 4;
2       ulong *x = new ulong[n];  // permutations
3       for (ulong k=0; k<n; ++k)  x[k] = k;
4       perm_heap2_swaps P(n);
5       do
6       {
7           ulong sw1, sw2;
8           P.get_swap(sw1, sw2);
9           swap2( x[sw1], x[sw2] );  // update permutation
10          // visit permutation
11      }
12      while ( P.next() );
```

The update routine works at a rate about 310 million per second (about 7 CPU cycles per update). Using arrays instead of pointers results in a rate of about 420 million per second (5.25 cycles per update).

Heap's algorithm and the optimization idea was taken from the excellent survey [235] which gives several permutation algorithms and implementations in pseudo code.

10.6 Lipski's Minimal-change orders

10.6.1 Variants of Heap's algorithm

Various algorithms similar to Heap's method are given in Lipski's paper [187], we take three of those and add a similar one. The four orderings obtained for the permutations of five elements are shown in figure 10.6-A. The leftmost order is Heap's order. The implementation is given in [FXT: `class perm_gray_lipski` in comb/perm-gray-lipski.h], the variable `r` determines the order that is generated:

	x=(j&1 ? 0 : d);		x=(j&1 ? 0 : j-d);		x=(j&1 ? j-1 : d);		x=(j&1 ? j-1 : j-d);	
1:	[. 1 2 3 4]		[. 1 2 3 4]		[. 1 2 3 4]		[. 1 2 3 4]	
2:	[1 . 2 3 4]	(1)	[1 . 2 3 4]	(1)	[1 . 2 3 4]	(1)	[1 . 2 3 4]	(1)
3:	[2 . 1 3 4]	(2)	[2 . 1 3 4]	(2)	[2 . 1 3 4]	(2)	[2 . 1 3 4]	(2)
4:	[. 2 1 3 4]	(1)	[. 2 1 3 4]	(1)	[. 2 1 3 4]	(1)	[. 2 1 3 4]	(1)
5:	[1 2 . 3 4]	(2)	[1 2 . 3 4]	(2)	[1 2 . 3 4]	(2)	[1 2 . 3 4]	(2)
6:	[2 1 . 3 4]	(1)	[2 1 . 3 4]	(1)	[2 1 . 3 4]	(1)	[2 1 . 3 4]	(1)
7:	[3 1 . 2 4]	(3)	[2 1 3 . 4]	(3,2)	[3 1 . 2 4]	(3)	[2 1 3 . 4]	(3,2)
8:	[1 3 . 2 4]	(1)	[1 2 3 . 4]	(1)	[1 3 . 2 4]	(1)	[1 2 3 . 4]	(1)
9:	[. 3 1 2 4]	(2)	[3 2 1 . 4]	(2)	[. 3 1 2 4]	(2)	[3 2 1 . 4]	(2)
10:	[3 . 1 2 4]	(1)	[2 3 1 . 4]	(1)	[3 . 1 2 4]	(1)	[2 3 1 . 4]	(1)
11:	[1 . 3 2 4]	(2)	[1 3 2 . 4]	(2)	[1 . 3 2 4]	(2)	[1 3 2 . 4]	(2)
12:	[. 1 3 2 4]	(1)	[3 1 2 . 4]	(1)	[. 1 3 2 4]	(1)	[3 1 2 . 4]	(1)
13:	[. 2 3 1 4]	(3,1)	[3 . 2 1 4]	(3,1)	[. 2 3 1 4]	(3,1)	[3 . 2 1 4]	(3,1)
14:	[2 . 3 1 4]	(1)	[. 3 2 1 4]	(1)	[2 . 3 1 4]	(1)	[. 3 2 1 4]	(1)
15:	[3 . 2 1 4]	(2)	[2 3 . 1 4]	(2)	[3 . 2 1 4]	(2)	[2 3 . 1 4]	(2)
16:	[. 3 2 1 4]	(1)	[3 2 . 1 4]	(1)	[. 3 2 1 4]	(1)	[3 2 . 1 4]	(1)
17:	[2 3 . 1 4]	(2)	[. 2 3 1 4]	(2)	[2 3 . 1 4]	(2)	[. 2 3 1 4]	(2)
18:	[3 2 . 1 4]	(1)	[2 . 3 1 4]	(1)	[3 2 . 1 4]	(1)	[2 . 3 1 4]	(1)
19:	[3 2 1 . 4]	(3,2)	[1 . 3 2 4]	(3)	[3 2 1 . 4]	(3,2)	[1 . 3 2 4]	(3)
20:	[2 3 1 . 4]	(1)	[. 1 3 2 4]	(1)	[2 3 1 . 4]	(1)	[. 1 3 2 4]	(1)
21:	[1 3 2 . 4]	(2)	[3 1 . 2 4]	(2)	[1 3 2 . 4]	(2)	[3 1 . 2 4]	(2)
22:	[3 1 2 . 4]	(1)	[1 3 . 2 4]	(1)	[3 1 2 . 4]	(1)	[1 3 . 2 4]	(1)
23:	[2 1 3 . 4]	(2)	[. 3 1 2 4]	(2)	[2 1 3 . 4]	(2)	[. 3 1 2 4]	(2)
24:	[1 2 3 . 4]	(1)	[3 . 1 2 4]	(1)	[1 2 3 . 4]	(1)	[3 . 1 2 4]	(1)
25:	[4 2 3 . 1]	(4)	[4 . 1 2 3]	(4)	[1 2 4 . 3]	(4,2)	[3 . 4 2 1]	(4,2)
26:	[2 4 3 . 1]	(1)	[. 4 1 2 3]	(1)	[2 1 4 . 3]	(1)	[. 3 4 2 1]	(1)
27:	[3 4 2 . 1]	(2)	[1 4 . 2 3]	(2)	[4 1 2 . 3]	(2)	[4 3 . 2 1]	(2)
28:	[4 3 2 . 1]	(1)	[4 1 . 2 3]	(1)	[1 4 2 . 3]	(1)	[3 4 . 2 1]	(1)
29:	[2 3 4 . 1]	(2)	[. 1 4 2 3]	(2)	[2 4 1 . 3]	(2)	[. 4 3 2 1]	(2)
30:	[3 2 4 . 1]	(1)	[1 . 4 2 3]	(1)	[4 2 1 . 3]	(1)	[4 . 3 2 1]	(1)
31:	[. 2 4 3 1]	(3)	[1 . 2 4 3]	(3,2)	[. 2 1 4 3]	(3)	[4 . 2 3 1]	(3,2)
32:	[2 . 4 3 1]	(1)	[. 1 2 4 3]	(1)	[2 . 1 4 3]	(1)	[. 4 2 3 1]	(1)
33:	[4 . 2 3 1]	(2)	[2 1 . 4 3]	(2)	[1 . 2 4 3]	(2)	[2 4 . 3 1]	(2)
34:	[. 4 2 3 1]	(1)	[1 2 . 4 3]	(1)	[. 1 2 4 3]	(1)	[4 2 . 3 1]	(1)
35:	[2 4 . 3 1]	(2)	[. 2 1 4 3]	(2)	[2 1 . 4 3]	(2)	[. 2 4 3 1]	(2)
36:	[4 2 . 3 1]	(1)	[2 . 1 4 3]	(1)	[1 2 . 4 3]	(1)	[2 . 4 3 1]	(1)
37:	[4 3 . 2 1]	(3,1)	[2 4 1 . 3]	(3,1)	[1 4 . 2 3]	(3,1)	[2 3 4 . 1]	(3,1)
38:	[3 4 . 2 1]	(1)	[4 2 1 . 3]	(1)	[4 1 . 2 3]	(1)	[3 2 4 . 1]	(1)
39:	[. 4 3 2 1]	(2)	[1 2 4 . 3]	(2)	[. 1 4 2 3]	(2)	[4 2 3 . 1]	(2)
40:	[4 . 3 2 1]	(1)	[2 1 4 . 3]	(1)	[1 . 4 2 3]	(1)	[2 4 3 . 1]	(1)
41:	[3 . 4 2 1]	(2)	[4 1 2 . 3]	(2)	[4 . 1 2 3]	(2)	[3 4 2 . 1]	(2)
42:	[. 3 4 2 1]	(1)	[1 4 2 . 3]	(1)	[. 4 1 2 3]	(1)	[4 3 2 . 1]	(1)
43:	[. 3 2 4 1]	(3,2)	[. 4 2 1 3]	(3)	[. 4 2 1 3]	(3,2)	[. 3 2 4 1]	(3)
44:	[3 . 2 4 1]	(1)	[4 . 2 1 3]	(1)	[4 . 2 1 3]	(1)	[3 . 2 4 1]	(1)
45:	[2 . 3 4 1]	(2)	[2 . 4 1 3]	(2)	[2 . 4 1 3]	(2)	[2 . 3 4 1]	(2)
46:	[. 2 3 4 1]	(1)	[. 2 4 1 3]	(1)	[. 2 4 1 3]	(1)	[. 2 3 4 1]	(1)
47:	[3 2 . 4 1]	(2)	[4 2 . 1 3]	(2)	[4 2 . 1 3]	(2)	[3 2 . 4 1]	(2)
48:	[2 3 . 4 1]	(1)	[2 4 . 1 3]	(1)	[2 4 . 1 3]	(1)	[2 3 . 4 1]	(1)
49:	[1 3 . 4 2]	(4)	[3 4 . 1 2]	(4)	[2 4 3 1 .]	(4,2)	[2 3 1 4 .]	(4,2)
50:	[3 1 . 4 2]	(1)	[4 3 . 1 2]	(1)	[4 2 3 1 .]	(1)	[3 2 1 4 .]	(1)
51:	[. 1 3 4 2]	(2)	[. 3 4 1 2]	(2)	[3 2 4 1 .]	(2)	[1 2 3 4 .]	(2)
52:	[1 . 3 4 2]	(1)	[3 . 4 1 2]	(1)	[2 3 4 1 .]	(1)	[2 1 3 4 .]	(1)
53:	[3 . 1 4 2]	(2)	[4 . 3 1 2]	(2)	[4 3 2 1 .]	(2)	[3 1 2 4 .]	(2)
54:	[. 3 1 4 2]	(1)	[. 4 3 1 2]	(1)	[3 4 2 1 .]	(1)	[1 3 2 4 .]	(1)
55:	[4 3 1 . 2]	(3)	[. 4 1 3 2]	(3,2)	[1 4 2 3 .]	(3)	[1 3 4 2 .]	(3,2)
56:	[3 4 1 . 2]	(1)	[4 . 1 3 2]	(1)	[4 1 2 3 .]	(1)	[3 1 4 2 .]	(1)
57:	[1 4 3 . 2]	(2)	[1 . 4 3 2]	(2)	[2 1 4 3 .]	(2)	[4 1 3 2 .]	(2)
58:	[4 1 3 . 2]	(1)	[. 1 4 3 2]	(1)	[1 2 4 3 .]	(1)	[1 4 3 2 .]	(1)
59:	[3 1 4 . 2]	(2)	[4 1 . 3 2]	(2)	[4 2 1 3 .]	(2)	[3 4 1 2 .]	(2)
60:	[1 3 4 . 2]	(1)	[1 4 . 3 2]	(1)	[2 4 1 3 .]	(1)	[4 3 1 2 .]	(1)
108:	[3 4 2 1 .]	(1)	[4 2 3 1 .]	(1)	[3 . 4 1 2]	(1)	[. 4 3 1 2]	(1)
109:	[3 1 2 4 .]	(3,1)	[4 1 3 2 .]	(3,1)	[3 1 4 . 2]	(3,1)	[. 1 3 4 2]	(3,1)
110:	[1 3 2 4 .]	(1)	[1 4 3 2 .]	(1)	[1 3 4 . 2]	(1)	[1 . 3 4 2]	(1)
111:	[2 3 1 4 .]	(2)	[3 4 1 2 .]	(2)	[4 3 1 . 2]	(2)	[3 . 1 4 2]	(2)
112:	[3 2 1 4 .]	(1)	[4 3 1 2 .]	(1)	[3 4 1 . 2]	(1)	[. 3 1 4 2]	(1)
113:	[1 2 3 4 .]	(2)	[1 3 4 2 .]	(2)	[1 4 3 . 2]	(2)	[1 3 . 4 2]	(2)
114:	[2 1 3 4 .]	(1)	[3 1 4 2 .]	(1)	[4 1 3 . 2]	(1)	[3 1 . 4 2]	(1)
115:	[2 1 4 3 .]	(3,2)	[2 1 4 3 .]	(3)	[4 1 . 3 2]	(3,2)	[4 1 . 3 2]	(3)
116:	[1 2 4 3 .]	(1)	[1 2 4 3 .]	(1)	[1 4 . 3 2]	(1)	[1 4 . 3 2]	(1)
117:	[4 2 1 3 .]	(2)	[4 2 1 3 .]	(2)	[. 4 1 3 2]	(2)	[. 4 1 3 2]	(2)
118:	[2 4 1 3 .]	(1)	[2 4 1 3 .]	(1)	[4 . 1 3 2]	(1)	[4 . 1 3 2]	(1)
119:	[1 4 2 3 .]	(2)	[1 4 2 3 .]	(2)	[1 . 4 3 2]	(2)	[1 . 4 3 2]	(2)
120:	[4 1 2 3 .]	(1)	[4 1 2 3 .]	(1)	[. 1 4 3 2]	(1)	[. 1 4 3 2]	(1)

Figure 10.6-A: First half and last permutations of five elements as obtained by variants of Heap's method. Next to the permutations the swaps are shown as (x, y), a swap $(x, 0)$ is given as (x).

```
1   class perm_gray_lipski
2   {
3     [--snip--]
4       ulong r_;   // order (0<=r<4):
5     [--snip--]
6
7       bool next()
8       {
9           // increment mixed radix number:
10          ulong j = 0;
11          while ( d_[j]==j+1 )  { d_[j]=0; ++j; }
12          if ( j<n_-1 ) // only if no overflow
13          {
14              const ulong d = d_[j];
15
16              ulong x;
17              switch ( r_ )
18              {
19              case 0:  x = (j&1 ? 0 : d);  break;     // Lipski(9) == Heap
20              case 1:  x = (j&1 ? 0 : j-d);  break;   // Lipski(16)
21              case 2:  x = (j&1 ? j-1 : d);  break;   // Lipski(10)
22              default: x = (j&1 ? j-1 : j-d);  break; // not in Lipski's paper
23              }
24              const ulong k = j+1;
25              swap2(p_[k], p_[x]);
26              sw1_ = k;  sw2_ = x;
27
28              d_[j] = d + 1;
29              return true;
30          }
31          else  return  false;  // j==n-1 for last permutation
32      }
33    [--snip--]
34  };
```

The top lines in figure 10.6-A repeat the statements in the switch-block. For three or less elements all orderings coincide, with $n = 4$ elements the orderings for $r = 0$ and $r = 2$, and the orderings for $r = 1$ and $r = 3$ coincide. About 110 million permutations per second are generated [FXT: comb/perm-gray-lipski-demo.cc]. Optimizations similar to those for Heaps method should be obvious.

10.6.2 Variants of Wells' algorithm

A Gray code for permutations given by Wells [269] is shown in the left of figure 10.6-B. The following implementation, following Lipski's paper [187], includes two variants of the algorithm. We just give the crucial assignments in the method that computes the successor [FXT: `class perm_gray_wells` in comb/perm-gray-wells.h]:

```
1   class perm_gray_wells
2   {
3     [--snip--]
4               switch ( r_ )
5               {
6               case 1:  x = ( (j&1) || (d==0) ? 0 : d-1);  break;  // Lipski(14)
7               case 2:  x = ( (j&1) || (d==0) ? j : d-1);  break;  // Lipski(15)
8               default: x = ( (j&1) || (d<=1) ? j : j-d);  break;  // Wells' order == Lipski(8)
9               }
10    [--snip--]
11  };
```

Both expressions `(d==0)` can be changed to `(d<=1)` without changing the algorithm. More than 90 million permutations per second are generated [FXT: comb/perm-gray-wells-demo.cc].

```
x=( (j&1) || (d<=1) ? j : j-d);          x=( (j&1) || (d==0) ? 0 : d-1);
   1:    [ . 1 2 3 ]                        1:    [ . 1 2 3 ]
   2:    [ 1 . 2 3 ]  (1, 0)                2:    [ 1 . 2 3 ]  (1, 0)
   3:    [ 1 2 . 3 ]  (2, 1)                3:    [ 2 . 1 3 ]  (2, 0)
   4:    [ 2 1 . 3 ]  (1, 0)                4:    [ . 2 1 3 ]  (1, 0)
   5:    [ 2 . 1 3 ]  (2, 1)                5:    [ 1 2 . 3 ]  (2, 0)
   6:    [ . 2 1 3 ]  (1, 0)                6:    [ 2 1 . 3 ]  (1, 0)
   7:    [ . 2 3 1 ]  (3, 2)                7:    [ 3 1 . 2 ]  (3, 0)
   8:    [ 2 . 3 1 ]  (1, 0)                8:    [ 1 3 . 2 ]  (1, 0)
   9:    [ 2 3 . 1 ]  (2, 1)                9:    [ . 3 1 2 ]  (2, 0)
  10:    [ 3 2 . 1 ]  (1, 0)               10:    [ 3 . 1 2 ]  (1, 0)
  11:    [ 3 . 2 1 ]  (2, 1)               11:    [ 1 . 3 2 ]  (2, 0)
  12:    [ . 3 2 1 ]  (1, 0)               12:    [ . 1 3 2 ]  (1, 0)
  13:    [ . 3 1 2 ]  (3, 2)               13:    [ 2 1 3 . ]  (3, 0)
  14:    [ 3 . 1 2 ]  (1, 0)               14:    [ 1 2 3 . ]  (1, 0)
  15:    [ 3 1 . 2 ]  (2, 1)               15:    [ 3 2 1 . ]  (2, 0)
  16:    [ 1 3 . 2 ]  (1, 0)               16:    [ 2 3 1 . ]  (1, 0)
  17:    [ 1 . 3 2 ]  (2, 1)               17:    [ 1 3 2 . ]  (2, 0)
  18:    [ . 1 3 2 ]  (1, 0)               18:    [ 3 1 2 . ]  (1, 0)
  19:    [ 2 1 3 . ]  (3, 0)               19:    [ 3 . 2 1 ]  (3, 1)
  20:    [ 1 2 3 . ]  (1, 0)               20:    [ . 3 2 1 ]  (1, 0)
  21:    [ 1 3 2 . ]  (2, 1)               21:    [ 2 3 . 1 ]  (2, 0)
  22:    [ 3 1 2 . ]  (1, 0)               22:    [ 3 2 . 1 ]  (1, 0)
  23:    [ 3 2 1 . ]  (2, 1)               23:    [ . 2 3 1 ]  (2, 0)
  24:    [ 2 3 1 . ]  (1, 0)               24:    [ 2 . 3 1 ]  (1, 0)
```

Figure 10.6-B: Wells' order for the permutations of four elements (left), and an order where most swaps are with the first position (right). Dots denote the element zero.

```
        permutation     swap      inverse p.     direction
 0:    [ . 1 2 3 ]    (3, 2)     [ . 1 2 3 ]     + + + +
 1:    [ 1 . 2 3 ]    (0, 1)     [ 1 . 2 3 ]     + + + +
 2:    [ 1 2 . 3 ]    (1, 2)     [ 2 . 1 3 ]     + + + +
 3:    [ 1 2 3 . ]    (2, 3)     [ 3 . 1 2 ]     + + + +
 4:    [ 2 1 3 . ]    (0, 1)     [ 3 1 . 2 ]     - + + +
 5:    [ 2 1 . 3 ]    (3, 2)     [ 2 1 . 3 ]     - + + +
 6:    [ 2 . 1 3 ]    (2, 1)     [ 1 2 . 3 ]     - + + +
 7:    [ . 2 1 3 ]    (1, 0)     [ . 2 1 3 ]     - + + +
 8:    [ . 2 3 1 ]    (2, 3)     [ . 3 1 2 ]     + + + +
 9:    [ 2 . 3 1 ]    (0, 1)     [ 1 3 . 2 ]     + + + +
10:    [ 2 3 . 1 ]    (1, 2)     [ 2 3 . 1 ]     + + + +
11:    [ 2 3 1 . ]    (2, 3)     [ 3 2 . 1 ]     + + + +
12:    [ 3 2 1 . ]    (0, 1)     [ 3 2 1 . ]     - - + +
13:    [ 3 2 . 1 ]    (3, 2)     [ 2 3 1 . ]     - - + +
14:    [ 3 . 2 1 ]    (2, 1)     [ 1 3 2 . ]     - - + +
15:    [ . 3 2 1 ]    (1, 0)     [ . 3 2 1 ]     - - + +
16:    [ . 3 1 2 ]    (3, 2)     [ . 2 3 1 ]     + - + +
17:    [ 3 . 1 2 ]    (0, 1)     [ 1 2 3 . ]     + - + +
18:    [ 3 1 . 2 ]    (1, 2)     [ 2 1 3 . ]     + - + +
19:    [ 3 1 2 . ]    (2, 3)     [ 3 1 2 . ]     + - + +
20:    [ 1 3 2 . ]    (1, 0)     [ 3 . 2 1 ]     - - + +
21:    [ 1 3 . 2 ]    (3, 2)     [ 2 . 3 1 ]     - - + +
22:    [ 1 . 3 2 ]    (2, 1)     [ 1 . 3 2 ]     - - + +
23:    [ . 1 3 2 ]    (1, 0)     [ . 1 3 2 ]     - - + +
```

Figure 10.7-A: The permutations of 4 elements in a strong minimal-change order (smallest element moves most often). Dots denote zeros.

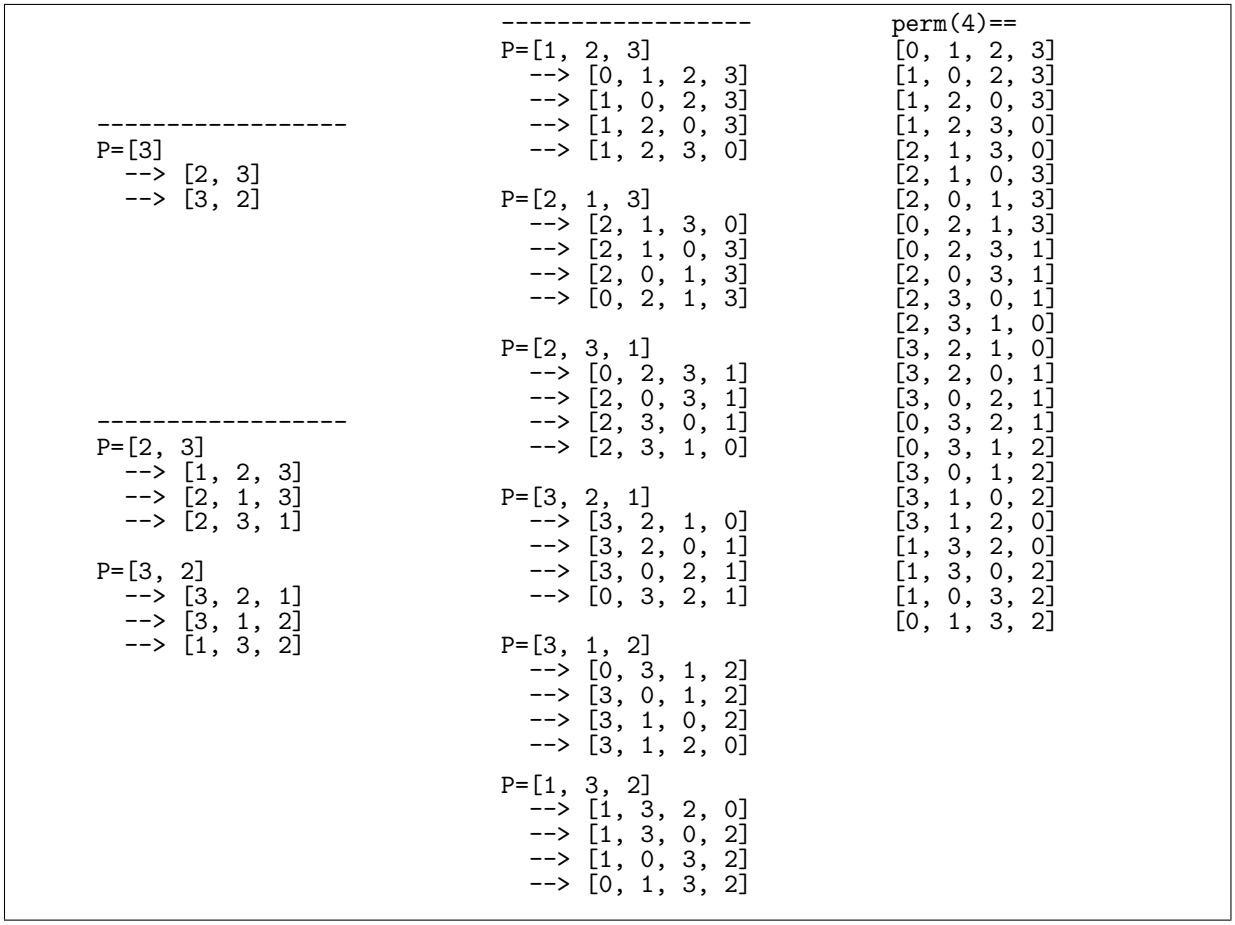

Figure 10.7-B: Trotter's construction as an interleaving process.

10.7 Strong minimal-change order (Trotter's algorithm)

10.7.1 Variant where smallest element moves most often

Figure 10.7-A shows the permutations of 4 elements in a *strong minimal-change order*: just two elements are swapped with each update and these are adjacent. Note that, in the sequence of the inverse permutations the swapped pair always consists of elements x and $x+1$. Also the first and last permutation differ by an adjacent transposition (of the last two elements). The ordering can be obtained by an interleaving process shown in figure 10.7-B. The first half of the permutations in this order are the reversals of the second half: the relative order of the two smallest elements is changed only with the transition just after the first half and reversal changes the order of these two elements. Mutual permutations lie $n!/2$ positions apart.

A computer program to obtain all permutations in the shown order was given 1962 by H. F. Trotter [256] (see also [154] and [113]). However, the order was already known long before in connection with bell ringing, under the name *Plain Changes*, see [170].

We compute both the permutation and its inverse [FXT: `class perm_trotter` in comb/perm-trotter.h]:

```
1   class perm_trotter
2   {
3   public:
4       ulong n_;     // number of elements to permute
5       ulong *x_;    // permutation of {0, 1, ..., n-1}
6       ulong *xi_;   // inverse permutation
```

```
7       ulong *d_;    // auxiliary: directions
8       ulong sw1_, sw2_; // indices of elements swapped most recently
9
10    public:
11      perm_trotter(ulong n)
12      {
13          n_ = n;
14          x_ = new ulong[n_+2];
15          xi_ = new ulong[n_];
16          d_ = new ulong[n_];
17          ulong sen = 0;  // sentinel value minimal
18          x_[0] = x_[n_+1] = sen;
19          ++x_;
20          first();
21      }
22     [--snip--]
23
```

Note that sentinel elements are at the lower and higher end of the array for the permutation. For each element we store a direction-flag $= \pm 1$ in an array `d_[]`. Initially all are set to $+1$:

```
1       void first()
2       {
3           for (ulong i=0; i<n_; i++)  xi_[i] = i;
4           for (ulong i=0; i<n_; i++)  x_[i] = i;
5           for (ulong i=0; i<n_; i++)  d_[i] = 1;
6           sw1_ = n_ - 1;  sw2_ = n_ - 2;  // relative to last permutation
7       }
8     [--snip--]
```

To compute the successor, find the smallest element `e1` whose neighbor `e2` (left or right neighbor, according to the direction) is greater than `e1`. Swap the elements `e1` with `e2` and change the direction of all elements that could not be moved. The location of the elements, `i1` and `i2` are found with the inverse permutation, which has to be updated accordingly:

```
1       bool next()
2       {
3           for (ulong e1=0; e1<n_; ++e1)
4           {
5               // e1 is the element we try to move
6               ulong i1 = xi_[e1];   // position of element e1
7               ulong d = d_[e1];     // direction to move e1
8               ulong i2 = i1 + d;    // position to swap with
9               ulong e2 = x_[i2];    // element to swap with
10
11              if ( e1 < e2 )  // can we swap?
12              {
13                  xi_[e1] = i2;
14                  xi_[e2] = i1;
15                  x_[i1] = e2;
16                  x_[i2] = e1;
17                  sw1_ = i1;  sw2_ = i2;
18                  while ( e1-- )  d_[e1] = -d_[e1];
19                  return true;
20              }
21          }
22
23          first();
24          return false;
25      }
```

The locations of the swap can be obtained with the method

```
1       void get_swap(ulong &s1, ulong &s2)  const
2       { s1=sw1_; s2=sw2_; }
```

The last permutation can be obtained via

```
1       void last()
2       {
3           for (ulong i=0; i<n_; i++)  xi_[i] = i;
4           for (ulong i=0; i<n_; i++)  x_[i] = i;
5           for (ulong i=0; i<n_; i++)  d_[i] = -1UL;
6           sw1_ = n_ - 1;  sw2_ = n_ - 2;  // relative to first permutation
7           d_[sw1_] = +1;
8           d_[sw2_] = +1;
```

```
9           swap2(x_[sw1_], x_[sw2_]);
10          swap2(xi_[sw1_], xi_[sw2_]);
11      }
```

The routine for the predecessor is obtained by adding *one character* to the routine `next()`, it's a minus:

```
1       bool prev()
2       {
3     [--snip--]
4               ulong d = -d_[e1];     // direction to move e1 (NOTE: negated)
5     [--snip--]
6           last();
7           return false;
8       }
```

Well, we also changed the call `first()` to `last()`.

The routines `next()` and `prev()` generate about 137 million permutations per second. Figure 10.7-A was created with the program [FXT: comb/perm-trotter-demo.cc]:

```
ulong n = 4;
perm_trotter P(n);
do
{
    // visit permutation
}
while ( P.next() );
```

10.7.2 Optimized routines

The element zero is moved most often, so we can treat that case separately [FXT: comb/perm-trotter.h]:

```
1       bool next()
2       {
3           {  // most frequent case: e1 == 0
4               ulong i1 = xi_[0];   // position of element e1
5               ulong d = d_[0];     // direction to move e1
6               ulong i2 = i1 + d;    // position to swap with
7               ulong e2 = x_[i2];    // element to swap with
8
9               if ( 0 < e2 )  // can we swap?
10              {
11                  xi_[0] = i2;
12                  xi_[e2] = i1;
13                  x_[i1] = e2;
14                  x_[i2] = 0;
15                  sw1_ = i1;  sw2_ = i2;
16                  return true;
17              }
18          }
19          for (ulong e1=1; e1<n_; ++e1)  // note: start at e1=1
20      [--snip--]
```

The very same modification can be applied to the method `prev()`, only the minus has to be added:

```
            ulong d = -d_[0];     // direction to move e1 (NOTE: negated)
```

Now both methods compute about 174 million permutations per second, corresponding to less than 12.6 CPU cycles per update.

We can also treat the second most frequent case separately by adding the block:

```
1       bool next()
2       {
3           {  // most frequent case: e1 == 0
4     [--snip--]
5           }
6           {  // second most frequent case: e1 == 1
7               ulong i1 = xi_[1];   // position of element e1
8               ulong d = d_[1];     // direction to move e1
9               ulong i2 = i1 + d;    // position to swap with
10              ulong e2 = x_[i2];    // element to swap with
11
```

```
12              if ( 1 < e2 )  // can we swap?
13              {
14                  xi_[1] = i2;
15                  xi_[e2] = i1;
16                  x_[i1] = e2;
17                  x_[i2] = 1;
18                  d_[0] = -d_[0];  // negate
19                  sw1_ = i1;  sw2_ = i2;
20                  return true;
21              }
22          }
23
24          for (ulong e1=2; e1<n_; ++e1)  // note: start at e1=2
25      [--snip--]
```

With this modification the rate increases to about 179 million per second (12.3 cycles per update).

10.7.3 Variant where largest element moves most often

	permutation	swap	inverse p.	direction
0:	[. 1 2 3]	(0, 1)	[. 1 2 3]	- - - -
1:	[. 1 3 2]	(3, 2)	[. 1 3 2]	- - - -
2:	[. 3 1 2]	(2, 1)	[. 2 3 1]	- - - -
3:	[3 . 1 2]	(1, 0)	[1 2 3 .]	- - - -
4:	[3 . 2 1]	(3, 2)	[1 3 2 .]	- - - +
5:	[. 3 2 1]	(0, 1)	[. 3 2 1]	- - - +
6:	[. 2 3 1]	(1, 2)	[. 3 1 2]	- - - +
7:	[. 2 1 3]	(2, 3)	[. 2 1 3]	- - - +
8:	[2 . 1 3]	(1, 0)	[1 2 . 3]	- - - -
9:	[2 . 3 1]	(3, 2)	[1 3 . 2]	- - - -
10:	[2 3 . 1]	(2, 1)	[2 3 . 1]	- - - -
11:	[3 2 . 1]	(1, 0)	[2 3 1 .]	- - - -
12:	[3 2 1 .]	(3, 2)	[3 2 1 .]	- - + +
13:	[2 3 1 .]	(0, 1)	[3 2 . 1]	- - + +
14:	[2 1 3 .]	(1, 2)	[3 1 . 2]	- - + +
15:	[2 1 . 3]	(2, 3)	[2 1 . 3]	- - + +
16:	[1 2 . 3]	(0, 1)	[2 . 1 3]	- - + -
17:	[1 2 3 .]	(3, 2)	[3 . 1 2]	- - + -
18:	[1 3 2 .]	(2, 1)	[3 . 2 1]	- - + -
19:	[3 1 2 .]	(1, 0)	[3 1 2 .]	- - + -
20:	[3 1 . 2]	(2, 3)	[2 1 3 .]	- - + +
21:	[1 3 . 2]	(0, 1)	[2 . 3 1]	- - + +
22:	[1 . 3 2]	(1, 2)	[1 . 3 2]	- - + +
23:	[1 . 2 3]	(2, 3)	[1 . 2 3]	- - + +

Figure 10.7-C: The permutations of 4 elements in a strong minimal-change order (largest element moves most often). Dots denote zeros.

A variant of the algorithm moves the largest element most often as shown in figure 10.7-C (created with [FXT: comb/perm-trotter-lg-demo.cc]). Only a few modifications have to be made to the code [FXT: `class perm_trotter_lg` in comb/perm-trotter-lg.h]. The sentinel needs to be greater than all elements of the permutations, and the directions start with minus one:

```
1   class perm_trotter_lg
2   {
3     [--snip--]
4   public:
5       perm_trotter_lg(ulong n)
6       {
7     [--snip--]
8           ulong sen = n_;  // sentinel value maximal
9           x_[0] = x_[n_+1] = sen;
10          ++x_;
11          first();
12      }
13    [--snip--]
14      void first()
```

```
15      {
16          for (ulong i=0; i<n_; i++)  xi_[i] = i;
17          for (ulong i=0; i<n_; i++)  x_[i] = i;
18          for (ulong i=0; i<n_; i++)  d_[i] = -1UL;
19          sw1_ = 0;  sw2_ = 1;  // relative to last permutation
20      }
21    [--snip--]
```

In the update routine we look for the largest element whose neighbor is smaller than itself:

```
1       bool next()
2       {
3           ulong e1 = n_;
4           while ( e1-- )
5           {
6               // e1 is the element we try to move
7               ulong i1 = xi_[e1];   // position of element e1
8               ulong d = d_[e1];     // direction to move e1
9               ulong i2 = i1 + d;    // position to swap with
10              ulong e2 = x_[i2];    // element to swap with
11
12              if ( e1 > e2 )  // can we swap?
13              {
14                  xi_[e1] = i2;
15                  xi_[e2] = i1;
16                  x_[i1] = e2;
17                  x_[i2] = e1;
18                  sw1_ = i1;  sw2_ = i2;
19                  while ( ++e1<n_ )  d_[e1] = -d_[e1];
20                  return true;
21              }
22          }
23
24          first();
25          return false;
26      }
```

The last permutation can be obtained via

```
1       void last()
2       {
3           for (ulong i=0; i<n_; i++)  xi_[i] = i;
4           for (ulong i=0; i<n_; i++)  x_[i] = i;
5
6           for (ulong i=0; i<n_; i++)  d_[i] = +1;
7           sw1_ = 0;  sw2_ = 1;  // relative to first permutation
8           d_[sw1_] = -1UL;
9           d_[sw2_] = -1UL;
10          swap2(x_[sw1_], x_[sw2_]);
11          swap2(xi_[sw1_], xi_[sw2_]);
12      }
```

The method to compute the predecessor is obtained as before. The routines `next()` and `prev()` generate about 126 million permutations per second.

10.8 Minimal-change orders from factorial numbers

10.8.1 Permutations with falling factorial numbers

The Gray code for the mixed radix numbers with falling factorial base allows the computation of the permutations in minimal-change order in an elegant way. See figure 10.8-A which was created with the program [FXT: comb/perm-gray-ffact2-demo.cc]. The algorithm is implemented in [FXT: `class perm_gray_ffact2` in comb/perm-gray-ffact2.h]:

```
1   class perm_gray_ffact2
2   {
3   public:
4       mixedradix_gray2 *mrg_;  // loopless routine
5       ulong n_;   // number of elements to permute
6       ulong *x_;  // current permutation (of {0, 1, ..., n-1})
```

	permutation	ffact	pos	dir	inverse perm.
0:	[. 1 2 3]	[. . .]			[. 1 2 3]
1:	[1 . 2 3]	[1 . .]	0	+1	[1 . 2 3]
2:	[1 2 . 3]	[2 . .]	0	+1	[2 . 1 3]
3:	[1 2 3 .]	[3 . .]	0	+1	[3 . 1 2]
4:	[2 1 3 .]	[3 1 .]	1	+1	[3 1 . 2]
5:	[2 1 . 3]	[2 1 .]	0	-1	[2 1 . 3]
6:	[2 . 1 3]	[1 1 .]	0	-1	[1 2 . 3]
7:	[. 2 1 3]	[. 1 .]	0	-1	[. 2 1 3]
8:	[. 2 3 1]	[. 2 .]	1	+1	[. 3 1 2]
9:	[2 . 3 1]	[1 2 .]	0	+1	[1 3 . 2]
10:	[2 3 . 1]	[2 2 .]	0	+1	[2 3 . 1]
11:	[2 3 1 .]	[3 2 .]	0	+1	[3 2 . 1]
12:	[3 2 1 .]	[3 2 1]	2	+1	[3 2 1 .]
13:	[3 2 . 1]	[2 2 1]	0	-1	[2 3 1 .]
14:	[3 . 2 1]	[1 2 1]	0	-1	[1 3 2 .]
15:	[. 3 2 1]	[. 2 1]	0	-1	[. 3 2 1]
16:	[. 3 1 2]	[. 1 1]	1	-1	[. 2 3 1]
17:	[3 . 1 2]	[1 1 1]	0	+1	[1 2 3 .]
18:	[3 1 . 2]	[2 1 1]	0	+1	[2 1 3 .]
19:	[3 1 2 .]	[3 1 1]	0	+1	[3 1 2 .]
20:	[1 3 2 .]	[3 . 1]	1	-1	[3 . 2 1]
21:	[1 3 . 2]	[2 . 1]	0	-1	[2 . 3 1]
22:	[1 . 3 2]	[1 . 1]	0	-1	[1 . 3 2]
23:	[. 1 3 2]	[. . 1]	0	-1	[. 1 3 2]

Figure 10.8-A: Permutations in minimal-change order (left) and Gray code for mixed radix numbers with falling factorial base. The two rightmost columns give the place of change with the mixed radix numbers and its direction. Whenever digit p (=pos) changes by $d = \pm 1$ (=dir) in the mixed radix sequence the element p of the permutation is swapped with its right ($d = +1$) or left ($d = -1$) neighbor.

```
7       ulong *ix_;  // inverse permutation
8       ulong sw1_, sw2_; // indices of elements swapped most recently
9
10    public:
11      perm_gray_ffact2(ulong n)
12          : n_(n)
13      {
14          x_ = new ulong[n_];
15          ix_ = new ulong[n_];
16          mrg_ = new mixedradix_gray2(n_-1, 0);  // falling factorial base
17          first();
18      }
19
20     [--snip--]
21
22      void first()
23      {
24          mrg_->first();
25          for (ulong k=0; k<n_; ++k)  x_[k] = ix_[k] = k;
26          sw1_=n_-1;  sw2_=n_-2;
27      }
```

The crucial part is the computation of the successor:

```
1       bool next()
2       {
3           // Compute next mixed radix number in Gray code order:
4           if ( false == mrg_->next() )  { first(); return false; }
5           const ulong j = mrg_->pos();  // position of changed digit
6           const int d = mrg_->dir();    // direction of change
7
8           // swap:
9           const ulong x1 = j;        // element j
10          const ulong i1 = ix_[x1];  // position of j
11          const ulong i2 = i1 + d;   // neighbor
12          const ulong x2 = x_[i2];   // position of neighbor
13          x_[i1] = x2;   x_[i2] = x1;   // swap2(x_[i1], x_[i2]);
14          ix_[x1] = i2;  ix_[x2] = i1;  // swap2(ix_[x1], ix_[x2]);
15          sw1_=i1;  sw2_=i2;
16          return true;
```

```
17      }
```

The class uses the loopless algorithm for the computation of the mixed radix Gray code, so it is loopless itself. An alternative (CAT) algorithm is implemented in [FXT: `class perm_gray_ffact` in comb/perm-gray-ffact.h], we give just the routine for the successor:

```
1   private:
2       void swap(ulong j, ulong im)  // used with next() and prev()
3       {
4           const ulong x1 = j;        // element j
5           const ulong i1 = ix_[x1];  // position of j
6           const ulong i2 = i1 + im;  // neighbor
7           const ulong x2 = x_[i2];   // position of neighbor
8           x_[i1] = x2;   x_[i2] = x1;   // swap2(x_[i1], x_[i2]);
9           ix_[x1] = i2;  ix_[x2] = i1;  // swap2(ix_[x1], ix_[x2]);
10          sw1_=i1;  sw2_=i2;
11      }
12
13  public:
14      bool next()
15      {
16          ulong j = 0;
17          ulong m1 = n_ - 1;  // nine in falling factorial base
18          ulong ij;
19          while ( (ij=i_[j]) )
20          {
21              ulong im = i_[j];
22              ulong dj = d_[j] + im;
23              if ( dj>m1 )  // =^= if ( (dj>m1) || ((long)dj<0) )
24              {
25                  i_[j] = -ij;
26              }
27              else
28              {
29                  d_[j] = dj;
30                  swap(j, im);
31                  return true;
32              }
33
34              --m1;
35              ++j;
36          }
37          return false;
38      }
```

The routine for the predecessor (`prev()`) is obtained by replacing the statement `ulong im = i_[j];` by `ulong im = -i_[j];`. The loopless routine computes about 80 million permutations per second, the CAT version about 110 million per second [FXT: comb/perm-gray-ffact-demo.cc]. Both are slower than the implementations given in section 10.7 on page 254.

10.8.2 Permutations with rising factorial numbers

Figure 10.8-B shows a Gray code for permutations based on the Gray code for numbers in rising factorial base. The ordering coincides with Heap's algorithm (see section 10.5 on page 248) for up to four elements. A recursive construction for the order is shown in figure 10.8-C. The figure was created with the program [FXT: comb/perm-gray-rfact-demo.cc] (see also [FXT: comb/fact2perm-demo.cc]). A constant amortized time (CAT) algorithm for generating the permutations is [FXT: `class perm_gray_rfact` in comb/perm-gray-rfact.h]:

```
1   class perm_gray_rfact
2   {
3   public:
4       mixedradix_gray *M_;  // loopless routine
5       ulong n_;   // number of elements to permute
6       ulong *x_;  // current permutation (of {0, 1, ..., n-1})
7       ulong *ix_;  // inverse permutation
8       ulong sw1_, sw2_; // indices of elements swapped most recently
9
10  public:
11      perm_gray_rfact(ulong n)
```

```
        permutation          rfact      pos   dir       inverse perm.
 0:    [ . 1 2 3 ]        [ . . . ]                    [ . 1 2 3 ]
 1:    [ 1 . 2 3 ]        [ 1 . . ]     0     +1       [ 1 . 2 3 ]
 2:    [ 2 . 1 3 ]        [ 1 1 . ]     1     +1       [ 1 2 . 3 ]
 3:    [ . 2 1 3 ]        [ . 1 . ]     0     -1       [ . 2 1 3 ]
 4:    [ 1 2 . 3 ]        [ . 2 . ]     1     +1       [ 2 . 1 3 ]
 5:    [ 2 1 . 3 ]        [ 1 2 . ]     0     +1       [ 2 1 . 3 ]
 6:    [ 3 1 . 2 ]        [ 1 2 1 ]     2     +1       [ 2 1 3 . ]
 7:    [ 1 3 . 2 ]        [ . 2 1 ]     0     -1       [ 2 . 3 1 ]
 8:    [ . 3 1 2 ]        [ . 1 1 ]     1     -1       [ . 2 3 1 ]
 9:    [ 3 . 1 2 ]        [ 1 1 1 ]     0     +1       [ 1 2 3 . ]
10:    [ 1 . 3 2 ]        [ 1 . 1 ]     1     -1       [ 1 . 3 2 ]
11:    [ . 1 3 2 ]        [ . . 1 ]     0     -1       [ . 1 3 2 ]
12:    [ . 2 3 1 ]        [ . . 2 ]     2     +1       [ . 3 1 2 ]
13:    [ 2 . 3 1 ]        [ 1 . 2 ]     0     +1       [ 1 3 . 2 ]
14:    [ 3 . 2 1 ]        [ 1 1 2 ]     1     +1       [ 1 3 2 . ]
15:    [ . 3 2 1 ]        [ . 1 2 ]     0     -1       [ . 3 2 1 ]
16:    [ 2 3 . 1 ]        [ . 2 2 ]     1     +1       [ 2 3 . 1 ]
17:    [ 3 2 . 1 ]        [ 1 2 2 ]     0     +1       [ 2 3 1 . ]
18:    [ 3 2 1 . ]        [ 1 2 3 ]     2     +1       [ 3 2 1 . ]
19:    [ 2 3 1 . ]        [ . 2 3 ]     0     -1       [ 3 2 . 1 ]
20:    [ 1 3 2 . ]        [ . 1 3 ]     1     -1       [ 3 . 2 1 ]
21:    [ 3 1 2 . ]        [ 1 1 3 ]     0     +1       [ 3 1 2 . ]
22:    [ 2 1 3 . ]        [ 1 . 3 ]     1     -1       [ 3 1 . 2 ]
23:    [ 1 2 3 . ]        [ . . 3 ]     0     -1       [ 3 . 1 2 ]
```

Figure 10.8-B: Permutations in minimal-change order (left) and Gray code for mixed radix numbers with rising factorial base. For even n the first and last permutations are cyclic shifts by one of each other.

```
perm(2)=
  01
  10

append 2:
  01 2
  10 2

reverse and swap (2,1)
  20 1
  02 1

reverse and swap (1,0)
  12 0
  21 0

==> perm(3)
  012
  102
  201
  021
  120
  210
```

```
append 3:
  012 3
  102 3
  201 3
  021 3
  120 3
  210 3

reverse and swap (3,2):
  310 2
  130 2
  031 2
  301 2
  103 2
  013 2

reverse and swap (2,1):
  023 1
  203 1
  302 1
  032 1
  230 1
  320 1

reverse and swap (1,0):
  321 0
  231 0
  132 0
  312 0
  213 0
  123 0
```

```
==> perm(4):
  0123
  1023
  2013
  0213
  1203
  2103
  3102
  1302
  0312
  3012
  1032
  0132
  0231
  2031
  3021
  0321
  2301
  3201
  3210
  2310
  1320
  3120
  2130
  1230
```

Figure 10.8-C: Recursive construction of the permutations.

```
12          : n_(n)
13      {
14          x_ = new ulong[n_];
15          ix_ = new ulong[n_];
16          M_ = new mixedradix_gray(n_-1, 1);  // rising factorial base
17          first();
18      }
19    [--snip--]
20      void first()
21      {
22          M_->first();
23          for (ulong k=0; k<n_; ++k)  x_[k] = ix_[k] = k;
24          sw1_=n_-1;  sw2_=n_-2;
25      }
```

Let $j \geq 0$ be the position of the digit changed with incrementing the mixed radix number, and $d = \pm 1$ the increment or decrement of that digit. The swap to obtain the successor permutation swaps the element x_1 at position $j+1$ with the element x_2 where x_2 is lying to the left of x_1 and it is the greatest element smaller than x_1 for $d > 0$, and the smallest element greater than x_1 for $d < 0$:

```
1       bool next()
2       {
3           // Compute next mixed radix number in Gray code order:
4           if ( false == M_->next() )  { first(); return false; }
5           ulong j = M_->pos();  // position of changed digit
6
7           if ( j<=1 )  // easy cases: swap == (0,j+1)
8           {
9               const ulong i2 = j+1;  // i1 == 0
10              const ulong x1 = x_[0],  x2 = x_[i2];
11              x_[0] = x2;    x_[i2] = x1;  // swap2(x_[i1], x_[i2]);
12              ix_[x1] = i2;  ix_[x2] = 0;  // swap2(ix_[x1], ix_[x2]);
13              sw1_=0;  sw2_=i2;
14              return true;
15          }
16          else
17          {
18              ulong i1 = j+1,  i2 = i1;
19              ulong x1 = x_[i1],  x2;
20              int d = M_->dir();    // direction of change
21              if ( d>0 )
22              {
23                  x2 = 0;
24                  for (ulong t=0; t<i1; ++t)  // search maximal smaller element left
25                  {
26                      ulong xt = x_[t];
27                      if ( (xt < x1) && (xt >= x2) )  { i2=t; x2=xt; }
28                  }
29              }
30              else
31              {
32                  x2 = n_;
33                  for (ulong t=0; t<i1; ++t)  // search minimal larger element
34                  {
35                      ulong xt = x_[t];
36                      if ( (xt > x1) && (xt <= x2) )  { i2=t; x2=xt; }
37                  }
38              }
39
40              x_[i1] = x2;    x_[i2] = x1;  // swap2(x_[i1], x_[i2]);
41              ix_[x1] = i2;  ix_[x2] = i1;  // swap2(ix_[x1], ix_[x2]);
42
43              sw1_=i2;  sw2_=i1;
44              return true;
45          }
46      }
```

There is a slightly more efficient algorithm to compute the successor using the inverse permutations:

```
1       bool next()
2       {
3     [--snip--]  /* easy cases as before */
4           else
5           {
6               ulong i1 = j+1,  i2 = i1;
7               ulong x1 = x_[i1],  x2;
8               int d = M_->dir();    // direction of change
```

```
9               if ( d>0 )  // in the inverse permutation search first smaller element left:
10              {
11                  for (x2=x1-1;  ; --x2)  if ( (i2=ix_[x2]) < i1 )  break;
12              }
13              else  // in the inverse permutation search first smaller element right:
14              {
15                  for (x2=x1+1;  ; ++x2)  if ( (i2=ix_[x2]) < i1 )  break;
16              }
17      [--snip--]  /* swaps as before */
18          }
19      }
```

The method is chosen by defining `SUCC_BY_INV` in the file [FXT: comb/perm-gray-rfact.h]. About 68 million permutations per second are generated, about 58 million with the first method.

10.8.3 Permutations with permuted factorial numbers

```
             permutation    swap        xfact      pos dir      inv.perm.
      0:    [ . 1 2 3 4 ]              [ . . . . ]              [ . 1 2 3 4 ]
      1:    [ 1 . 2 3 4 ]  (0, 1)     [ 1 . . . ]   0   +1     [ 1 . 2 3 4 ]
      2:    [ 2 . 1 3 4 ]  (0, 2)     [ 1 1 . . ]   1   +1     [ 1 2 . 3 4 ]
      3:    [ . 2 1 3 4 ]  (0, 1)     [ . 1 . . ]   0   -1     [ . 2 1 3 4 ]
      4:    [ 1 2 . 3 4 ]  (0, 2)     [ . 2 . . ]   1   +1     [ 2 . 1 3 4 ]
      5:    [ 2 1 . 3 4 ]  (0, 1)     [ 1 2 . . ]   0   +1     [ 2 1 . 3 4 ]
      6:    [ 2 1 . 4 3 ]  (3, 4)     [ 1 2 1 . ]   2   +1     [ 2 1 . 4 3 ]
      7:    [ 1 2 . 4 3 ]  (0, 1)     [ . 2 1 . ]   0   -1     [ 2 . 1 4 3 ]
   [--snip--]
     91:    [ 3 4 2 1 . ]  (0, 1)     [ . 2 4 3 ]   0   -1     [ 4 3 2 . 1 ]
     92:    [ 2 4 3 1 . ]  (0, 2)     [ . 1 4 3 ]   1   -1     [ 4 3 . 2 1 ]
     93:    [ 4 2 3 1 . ]  (0, 1)     [ 1 1 4 3 ]   0   +1     [ 4 3 1 2 . ]
     94:    [ 3 2 4 1 . ]  (0, 2)     [ 1 . 4 3 ]   1   -1     [ 4 3 1 . 2 ]
     95:    [ 2 3 4 1 . ]  (0, 1)     [ . . 4 3 ]   0   -1     [ 4 3 . 1 2 ]
     96:    [ 2 3 4 . 1 ]  (3, 4)     [ . . 3 3 ]   2   -1     [ 3 4 . 1 2 ]
     97:    [ 3 2 4 . 1 ]  (0, 1)     [ 1 . 3 3 ]   0   +1     [ 3 4 1 . 2 ]
   [--snip--]
    106:    [ 3 1 4 . 2 ]  (0, 2)     [ 1 . 2 3 ]   1   -1     [ 3 1 4 . 2 ]
    107:    [ 1 3 4 . 2 ]  (0, 1)     [ . . 2 3 ]   0   -1     [ 3 . 4 1 2 ]
    108:    [ 1 2 4 . 3 ]  (1, 4)     [ . . 1 3 ]   2   -1     [ 3 . 1 4 2 ]
    109:    [ 2 1 4 . 3 ]  (0, 1)     [ 1 . 1 3 ]   0   +1     [ 3 1 . 4 2 ]
    110:    [ 4 1 2 . 3 ]  (0, 2)     [ 1 1 1 3 ]   1   +1     [ 3 1 2 4 . ]
    111:    [ 1 4 2 . 3 ]  (0, 1)     [ . 1 1 3 ]   0   -1     [ 3 . 2 4 1 ]
    112:    [ 2 4 1 . 3 ]  (0, 2)     [ . 2 1 3 ]   1   +1     [ 3 2 . 4 1 ]
    113:    [ 4 2 1 . 3 ]  (0, 1)     [ 1 2 1 3 ]   0   +1     [ 3 2 1 4 . ]
    114:    [ 3 2 1 . 4 ]  (0, 4)     [ 1 2 . 3 ]   2   -1     [ 3 2 1 . 4 ]
    115:    [ 2 3 1 . 4 ]  (0, 1)     [ . 2 . 3 ]   0   -1     [ 3 2 . 1 4 ]
    116:    [ 1 3 2 . 4 ]  (0, 2)     [ . 1 . 3 ]   1   -1     [ 3 . 2 1 4 ]
    117:    [ 3 1 2 . 4 ]  (0, 1)     [ 1 1 . 3 ]   0   +1     [ 3 1 2 . 4 ]
    118:    [ 2 1 3 . 4 ]  (0, 2)     [ 1 . . 3 ]   1   -1     [ 3 1 . 2 4 ]
    119:    [ 1 2 3 . 4 ]  (0, 1)     [ . . . 3 ]   0   -1     [ 3 . 1 2 4 ]
```

Figure 10.8-D: Permutations with mixed radix numbers with radix vector $[2, 3, 5, 4]$.

The rising and falling factorial numbers are special cases of factorial numbers with permuted digits. We give a method to compute the Gray code for permutations from the Gray code for permuted (falling) factorial numbers. A permutation of the radices determines how often a digit at any position is changed: the leftmost changes most often, the rightmost least often. The permutations corresponding to the mixed radix numbers with radix vector $[2, 3, 5, 4]$, the falling factorial last two radices swapped, is shown in figure 10.8-D [FXT: comb/perm-gray-rot1-demo.cc]. The desired property of this ordering is that the last permutation is as close to a cyclic shift by one of the first as possible. With even n the Gray code with the falling factorial basis the last permutation is a shift by one. With odd n no such Gray code exists: the total number of transpositions with any Gray code is odd for all $n > 1$, but the cyclic rotation by one corresponds to an even number of transpositions. The best we can get is that the first e elements where $e \leq n$ is the greatest possible even number. For example,

```
first               last
```

```
n=6:   [ 0 1 2 3 4 5 ]     [ 1 2 3 4 5 0 ]
n=7:   [ 0 1 2 3 4 5 6 ]   [ 1 2 3 4 5 0 6 ]
```

We use this ordering to show the general method [FXT: `class perm_gray_rot1` in comb/perm-gray-rot1.h]:

```
1   class perm_gray_rot1
2   {
3   public:
4       mixedradix_gray *M_;  // Gray code for factorial numbers
5       ulong n_;   // number of elements to permute
6       ulong *x_;  // current permutation (of {0, 1, ..., n-1})
7       ulong *ix_;  // inverse permutation
8       ulong sw1_, sw2_; // indices of elements swapped most recently
9
10  public:
11      perm_gray_rot1(ulong n)
12      // Must have: n>=1
13      {
14          n_ = (n ? n : 1);  // at least one
15          x_ = new ulong[n_];
16          ix_ = new ulong[n_];
17
18          M_ = new mixedradix_gray(n_-1, 1);  // rising factorial base
19
20          // apply permutation of radix vector with mixed radix number:
21          if ( (n_ >= 3) && (n & 1) )  // odd n>=3
22          {
23              ulong *m1 = M_->m1_;
24              swap2(m1[n_-2], m1[n_-3]);  // swap last two factorial nines
25          }
26
27          first();
28      }
29    [--snip--]
```

The permutation applied here can be replaced by any permutation, the following update routines will still work:

```
1       bool next()
2       {
3           // Compute next mixed radix number in Gray code order:
4           if ( false == M_->next() )  { first();  return false; }
5
6           const ulong j = M_->pos();  // position of changed digit
7           const ulong i1 = M_->m1_[j];  // valid for any permutation of factorial radices
8
9           const ulong x1 = x_[i1];
10          ulong  i2 = i1,  x2;
11          const int d = M_->dir();    // direction of change
12
13          if ( d>0 )  // in the inverse permutation search first smaller element left:
14          {
15              for (x2=x1-1;  ; --x2)  if ( (i2=ix_[x2]) < i1 )  break;
16          }
17          else  // in the inverse permutation search first smaller element right:
18          {
19              for (x2=x1+1;  ; ++x2)  if ( (i2=ix_[x2]) < i1 )  break;
20          }
21
22          x_[i1] = x2;    x_[i2] = x1;  // swap2(x_[i1], x_[i2]);
23          ix_[x1] = i2;  ix_[x2] = i1;  // swap2(ix_[x1], ix_[x2]);
24
25          sw1_=i2;  sw2_=i1;
26
27          return true;
28      }
29    [--snip--]
```

Note that instead of taking $j+1$ as the position of the element to move, we take the value of the nine at the position j. The special ordering obtained here can be used to construct a Gray code with the single track property, see section 10.10.2 on page 272.

10.9 Orders where the smallest element always moves right

10.9.1 A variant of Trotter's construction

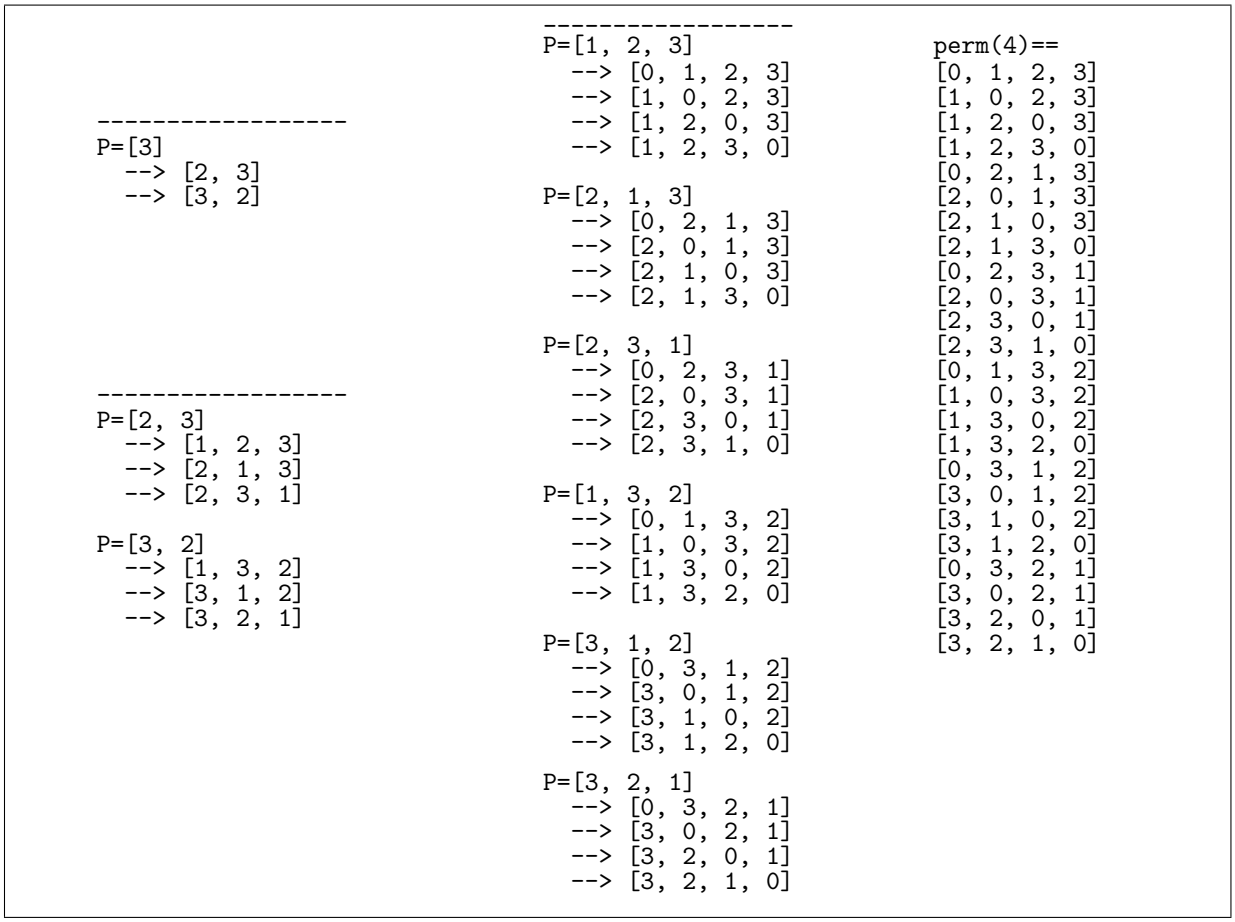

Figure 10.9-A: Interleaving process to obtain all permutations by right moves.

An ordering for the permutations where the first element always moves right can be obtained by the interleaving process shown in figure 10.9-A. The process is the same as that for Trotter's order shown in figure 10.7-B on page 254, but without changing the directions. The second half of the permutations is the reversed list of the reversed permutations in the first half. The permutations are shown in figure 10.9-B, they are the inverses of the permutations corresponding to the falling factorial numbers, see figure 10.3-A on page 238. An implementation is [FXT: `class perm_mv0` in comb/perm-mv0.h]:

```
1   class perm_mv0
2   {
3   public:
4       ulong *d_;  // mixed radix digits with radix = [n-1, n-2, n-3, ..., 2]
5       ulong *x_;  // permutation
6       ulong ect_; // counter for easy case
7       ulong n_;   // permutations of n elements
8
9   public:
10      perm_mv0(ulong n)
11      // Must have n>=2
12      {
13          n_ = n;
14          d_ = new ulong[n_];
15          d_[n-1] = 1;  // sentinel (must be nonzero)
16          x_ = new ulong[n_];
17          first();
18      }
```

```
              permutation     ffact       inv. perm.
        0:    [ . 1 2 3 ]    [ . . . ]    [ . 1 2 3 ]
        1:    [ 1 . 2 3 ]    [ 1 . . ]    [ 1 . 2 3 ]
        2:    [ 1 2 . 3 ]    [ 2 . . ]    [ 2 . 1 3 ]
        3:    [ 1 2 3 . ]    [ 3 . . ]    [ 3 . 1 2 ]
        4:    [ . 2 1 3 ]    [ . 1 . ]    [ . 2 1 3 ]
        5:    [ 2 . 1 3 ]    [ 1 1 . ]    [ 1 2 . 3 ]
        6:    [ 2 1 . 3 ]    [ 2 1 . ]    [ 2 1 . 3 ]
        7:    [ 2 1 3 . ]    [ 3 1 . ]    [ 3 1 . 2 ]
        8:    [ . 2 3 1 ]    [ . 2 . ]    [ . 3 1 2 ]
        9:    [ 2 . 3 1 ]    [ 1 2 . ]    [ 1 3 . 2 ]
       10:    [ 2 3 . 1 ]    [ 2 2 . ]    [ 2 3 . 1 ]
       11:    [ 2 3 1 . ]    [ 3 2 . ]    [ 3 2 . 1 ]
       12:    [ . 1 3 2 ]    [ . . 1 ]    [ . 1 3 2 ]
       13:    [ 1 . 3 2 ]    [ 1 . 1 ]    [ 1 . 3 2 ]
       14:    [ 1 3 . 2 ]    [ 2 . 1 ]    [ 2 . 3 1 ]
       15:    [ 1 3 2 . ]    [ 3 . 1 ]    [ 3 . 2 1 ]
       16:    [ . 3 1 2 ]    [ . 1 1 ]    [ . 2 3 1 ]
       17:    [ 3 . 1 2 ]    [ 1 1 1 ]    [ 1 2 3 . ]
       18:    [ 3 1 . 2 ]    [ 2 1 1 ]    [ 2 1 3 . ]
       19:    [ 3 1 2 . ]    [ 3 1 1 ]    [ 3 1 2 . ]
       20:    [ . 3 2 1 ]    [ . 2 1 ]    [ . 3 2 1 ]
       21:    [ 3 . 2 1 ]    [ 1 2 1 ]    [ 1 3 2 . ]
       22:    [ 3 2 . 1 ]    [ 2 2 1 ]    [ 2 3 1 . ]
       23:    [ 3 2 1 . ]    [ 3 2 1 ]    [ 3 2 1 . ]
```

Figure 10.9-B: All permutations of 4 elements and falling factorial numbers used to update the permutations. Dots denote zeros.

```
19      [--snip--]
20
21        void first()
22        {
23            for (ulong k=0; k<n_; ++k)  x_[k] = k;
24            for (ulong k=0; k<n_-1; ++k)  d_[k] = 0;
25            ect_ = 0;
26        }
27      [--snip--]
```

The update process uses the falling factorial numbers. Let j be the position where the digit is incremented, and d the value before the increment. The update

```
     permutation       ffact
                            v-- increment at j=2
    [ 4 2 3 5 1 0 ]    [ 5 4 1 1 . ]  <--= digit before increment is d=1
    [ 0 1 4 3 2 5 ]    [ . . 2 1 . ]
```

is done in three steps:

```
    [ 4 2 3 5 1 0 ]    [ 5 4 1 1 . ]
    [ 4 3 2 5 1 0 ]    [ 5 4 2 1 . ]   move element at position d=1 to the right
    [ * * 4 3 2 5 ]    [ * * 2 1 . ]   move all but j=2 elements to end
    [ 0 1 4 3 2 5 ]    [ . . 2 1 . ]   insert identical permutation at start
```

We treat the first digit separately as it changes most often (easy case):

```
1       bool next()
2       {
3           if ( ++ect_ < n_ ) // easy case
4           {
5               swap2(x_[ect_], x_[ect_-1]);
6               return true;
7           }
8           else
9           {
10              ect_ = 0;
11              ulong j = 1;
12              ulong m1 = n_ - 2;   // nine in falling factorial base
13              while ( d_[j]==m1 )  // find digit to increment
14              {
15                  d_[j] = 0;
16                  --m1;
17                  ++j;
```

```
18              }
19
20              if ( j==n_-1 )  return false;  // current permutation is last
21
22              const ulong dj = d_[j];
23              d_[j] = dj + 1;
24
25              // element at d[j] moves one position to the right:
26              swap2( x_[dj], x_[dj+1] );
27
28              { // move n-j elements to end:
29                  ulong s = n_-j,  d = n_;
30                  do
31                  {
32                      --s;
33                      --d;
34                      x_[d] = x_[s];
35                  }
36                  while ( s );
37              }
38
39              // fill in 0,1,2,..,j-1 at start:
40              for (ulong k=0; k<j; ++k)  x_[k] = k;
41
42              return true;
43          }
44      }
45  }
```

The routine generates about 160 million permutations per second [FXT: comb/perm-mv0-demo.cc].

10.9.2 Ives' algorithm

```
            permutation    inv. perm.
        1:  [ . 1 2 3 ]    [ . 1 2 3 ]
        2:  [ 1 . 2 3 ]    [ 1 . 2 3 ]
        3:  [ 1 2 . 3 ]    [ 2 . 1 3 ]
        4:  [ 1 2 3 . ]    [ 3 . 1 2 ]
        5:  [ . 2 3 1 ]    [ . 3 1 2 ]
        6:  [ 2 . 3 1 ]    [ 1 3 . 2 ]
        7:  [ 2 3 . 1 ]    [ 2 3 . 1 ]
        8:  [ 2 3 1 . ]    [ 3 2 . 1 ]
        9:  [ . 3 1 2 ]    [ . 2 3 1 ]
       10:  [ 3 . 1 2 ]    [ 1 2 3 . ]
       11:  [ 3 1 . 2 ]    [ 2 1 3 . ]
       12:  [ 3 1 2 . ]    [ 3 1 2 . ]  << only update with more
       13:  [ . 2 1 3 ]    [ . 2 1 3 ]  << than one transposition
       14:  [ 2 . 1 3 ]    [ 1 2 . 3 ]
       15:  [ 2 1 . 3 ]    [ 2 1 . 3 ]
       16:  [ 2 1 3 . ]    [ 3 1 . 2 ]
       17:  [ . 1 3 2 ]    [ . 1 3 2 ]
       18:  [ 1 . 3 2 ]    [ 1 . 3 2 ]
       19:  [ 1 3 . 2 ]    [ 2 . 3 1 ]
       20:  [ 1 3 2 . ]    [ 3 . 2 1 ]
       21:  [ . 3 2 1 ]    [ . 3 2 1 ]
       22:  [ 3 . 2 1 ]    [ 1 3 2 . ]
       23:  [ 3 2 . 1 ]    [ 2 3 1 . ]
       24:  [ 3 2 1 . ]    [ 3 2 1 . ]
```

Figure 10.9-C: All permutations of 4 elements in an order by Ives.

An ordering where most of the moves are a move by one to the right of the smallest element is shown in figure 10.9-C. With n elements only one in $n\,(n-1)$ moves is more than a transposition (only the update from 12 to 13 in figure 10.9-C). The second half of the list of permutations is the reversed list of the reversed permutations in the first half. The algorithm, given by Ives [151], is implemented in [FXT: `class perm_ives` in comb/perm-ives.h]:

```
1   class perm_ives
2   {
3   public:
```

```
4       ulong *p_;   // permutation
5       ulong *ip_;  // inverse permutation
6       ulong n_;    // permutations of n elements
7
8   public:
9       perm_ives(ulong n)
10      // Must have: n >= 2
11      {
12          n_ = n;
13          p_ = new ulong[n_];
14          ip_ = new ulong[n_];
15          first();
16      }
17    [--snip--]
```

The computation of the successor is

```
1       bool next()
2       {
3           ulong e1 = 0,  u = n_ - 1;
4           do
5           {
6               const ulong i1 = ip_[e1];
7               const ulong i2 = (i1==u ? e1 : i1+1 );
8               const ulong e2 = p_[i2];
9               p_[i1] = e2;  p_[i2] = e1;
10              ip_[e1] = i2;  ip_[e2] = i1;
11
12              if ( (p_[e1]!=e1) || (p_[u]!=u) )  return true;
13
14              ++e1;
15              --u;
16          }
17          while ( u > e1 );
18
19          return false;
20      }
21    [--snip--]
```

The rate of generation is about 180 M/s [FXT: comb/perm-ives-demo.cc]. Using arrays instead of pointers increases the rate to about 190 M/s.

As the easy case with the update (when just the first element is moved) occurs so often it is natural to create an extra branch for it. When the define for `PERM_IVES_OPT` is made before the class definition then a counter is created:

```
1   class perm_ives
2   {
3     [--snip--]
4       ulong ctm_;   // aux: counter for easy case
5       ulong ctm0_;  // aux: start value of ctm == n*(n-1)-1
6     [--snip--]
```

When the counter is nonzero the following update can be used:

```
1       bool next()
2       {
3           if ( ctm_-- )  // easy case
4           {
5               const ulong i1 = ip_[0];  // e1 == 0
6               const ulong i2 = (i1==n_-1 ? 0 : i1+1);
7               const ulong e2 = p_[i2];
8               p_[i1] = e2;  p_[i2] = 0;
9               ip_[0] = i2;  ip_[e2] = i1;
10              return true;
11          }
12          ctm_ = ctm0_;
13
14      [--snip--]   // rest as before
15      }
```

When arrays are used a minimal speedup is obtained (rate 192 M/s), when pointers are used, the effect is a notable slowdown (rate 163 M/s).

The greatest speedup is obtained by a modified condition in the loop:

```
if ( (p_[e1]^e1) | (p_[u]^u) )  return true;
// same as:   if ( (p_[e1]!=e1) || (p_[u]!=u) )  return true;
```

The rate is increased to almost 194 M/s. This optimization is activated by defining `PERM_IVES_OPT2`.

10.10 Single track orders

```
              permutation                 inv. perm.
 0:    [ . 2 3 1 ]    [ . . . ]    [ . 3 1 2 ]
 1:    [ . 3 2 1 ]    [ 1 . . ]    [ . 3 2 1 ]
 2:    [ . 3 1 2 ]    [ . 1 . ]    [ . 2 3 1 ]
 3:    [ . 2 1 3 ]    [ 1 1 . ]    [ . 2 1 3 ]
 4:    [ . 1 2 3 ]    [ . 2 . ]    [ . 1 2 3 ]
 5:    [ . 1 3 2 ]    [ 1 2 . ]    [ . 1 3 2 ]
 6:    [ 1 . 2 3 ]    [ . . 1 ]    [ 1 . 2 3 ]
 7:    [ 1 . 3 2 ]    [ 1 . 1 ]    [ 1 . 3 2 ]
 8:    [ 2 . 3 1 ]    [ . 1 1 ]    [ 1 3 . 2 ]
 9:    [ 3 . 2 1 ]    [ 1 1 1 ]    [ 1 3 2 . ]
10:    [ 3 . 1 2 ]    [ . 2 1 ]    [ 1 2 3 . ]
11:    [ 2 . 1 3 ]    [ 1 2 1 ]    [ 1 2 . 3 ]
12:    [ 3 1 . 2 ]    [ . . 2 ]    [ 2 1 3 . ]
13:    [ 2 1 . 3 ]    [ 1 . 2 ]    [ 2 1 . 3 ]
14:    [ 1 2 . 3 ]    [ . 1 2 ]    [ 2 . 1 3 ]
15:    [ 1 3 . 2 ]    [ 1 1 2 ]    [ 2 . 3 1 ]
16:    [ 2 3 . 1 ]    [ . 2 2 ]    [ 2 3 . 1 ]
17:    [ 3 2 . 1 ]    [ 1 2 2 ]    [ 2 3 1 . ]
18:    [ 2 3 1 . ]    [ . . 3 ]    [ 3 2 . 1 ]
19:    [ 3 2 1 . ]    [ 1 . 3 ]    [ 3 2 1 . ]
20:    [ 3 1 2 . ]    [ . 1 3 ]    [ 3 1 2 . ]
21:    [ 2 1 3 . ]    [ 1 1 3 ]    [ 3 1 . 2 ]
22:    [ 1 2 3 . ]    [ . 2 3 ]    [ 3 . 1 2 ]
23:    [ 1 3 2 . ]    [ 1 2 3 ]    [ 3 . 2 1 ]
```

Figure 10.10-A: Permutations of 4 elements in single track order. Dots denote zeros.

```
23  <--=  permutations of 2 elements
32

11 23 32  <--= concatenate rows and prepend new element

112332   <--= shift  0
321123   <--= shift  2
233211   <--= shift  4


000000 112332 321123 233211  <--= concatenate rows and prepend new element

000000 112332 321123 233211  <--= shift  0
233211 000000 112332 321123  <--= shift  6
321123 233211 000000 112332  <--= shift 12
112332 321123 233211 000000  <--= shift 18
```

Figure 10.10-B: Construction of the single track order for permutations of 4 elements.

Figure 10.10-A shows a *single track order* for the permutations of four elements. Each column in the list of permutations is a cyclic shift of the first column. A recursive construction for the ordering is shown in figure 10.10-B. The figure was created with the program [FXT: comb/perm-st-demo.cc] which uses [FXT: `class perm_st` in comb/perm-st.h]:

```
1   class perm_st
2   {
3   public:
4       ulong *d_;  // mixed radix digits with radix = [2, 3, 4, ..., n-1, (sentinel=-1)]
5       ulong *p_;  // permutation
6       ulong *pi_; // inverse permutation
7       ulong n_;   // permutations of n elements
8
9   public:
```

```
10      perm_st(ulong n)
11      {
12          n_ = n;
13          d_ = new ulong[n_];
14          p_ = new ulong[n_];
15          pi_ = new ulong[n_];
16          d_[n-1] = -1UL;  // sentinel
17          first();
18      }
19    [--snip--]
```

The first permutation is in enup order (see section 6.5.1 on page 185):

```
1       const ulong *data()  const { return p_; }
2       const ulong *invdata()  const  { return pi_; }
3
4       void first()
5       {
6           for (ulong k=0; k<n_-1; ++k)  d_[k] = 0;
7           for (ulong k=0, e=0;  k<n_;  ++k)
8           {
9               p_[k] = e;
10              pi_[e] = k;
11              e = next_enup(e, n_-1);
12          }
13      }
14    [--snip--]
```

The swap with the inverse permutations are determined by the rightmost position j changing with mixed radix counting with rising factorial base. We write -1 for the last element, -2 for the second last, and so on:

```
j   swaps
0:  (-2,-1)
1:  (-3,-2)
2:  (-4,-3) (-2,-1)
3:  (-5,-4) (-3,-2)
4:  (-6,-5) (-4,-3) (-2,-1)
5:  (-7,-6) (-5,-4) (-3,-2)
j:  (-j-2, -j-1) ... (-2-(j%1), -1-(j%1))
```

The computation of the successor is CAT:

```
1       bool next()
2       {
3           // increment mixed radix number:
4           ulong j = 0;
5           while ( d_[j]==j+1 )  { d_[j]=0;  ++j; }
6
7           if ( j==n_-1 )  return false; // current permutation is last
8           ++d_[j];
9
10          for (ulong e1=n_-2-j, e2=e1+1;  e2<n_;  e1+=2, e2+=2)
11          {
12              const ulong i1 = pi_[e1];  // position of element e1
13              const ulong i2 = pi_[e2];  // position of element e2
14              pi_[e1] = i2;
15              pi_[e2] = i1;
16              p_[i1] = e2;
17              p_[i2] = e1;
18          }
19
20          return true;
21      }
```

All swaps with the inverse permutations are of adjacent pairs. The reversals of the first half of all permutations lie in the second half, the reversal of the k-the permutation lies at position $n! - 1 - k$

The single track property is independent of the first permutation, one can start with the trivial permutation using

```
1       void first_id()  // start with identical permutation
2       {
3           for (ulong k=0; k<n_-1; ++k)  d_[k] = 0;
4           for (ulong k=0; k<n_; ++k)  p_[k] = pi_[k] = k;
5       }
```

```
          permutation                 inv. perm.
 0:    [ . 1 2 3 ]    [ . . . ]    [ . 1 2 3 ]
 1:    [ . 1 3 2 ]    [ 1 . . ]    [ . 1 3 2 ]
 2:    [ . 2 3 1 ]    [ . 1 . ]    [ . 3 1 2 ]
 3:    [ . 3 2 1 ]    [ 1 1 . ]    [ . 3 2 1 ]
 4:    [ . 3 1 2 ]    [ . 2 . ]    [ . 2 3 1 ]
 5:    [ . 2 1 3 ]    [ 1 2 . ]    [ . 2 1 3 ]
 6:    [ 1 3 . 2 ]    [ . . 1 ]    [ 2 . 3 1 ]
 7:    [ 1 2 . 3 ]    [ 1 . 1 ]    [ 2 . 1 3 ]
 8:    [ 2 1 . 3 ]    [ . 1 1 ]    [ 2 1 . 3 ]
 9:    [ 3 1 . 2 ]    [ 1 1 1 ]    [ 2 1 3 . ]
10:    [ 3 2 . 1 ]    [ . 2 1 ]    [ 2 3 1 . ]
11:    [ 2 3 . 1 ]    [ 1 2 1 ]    [ 2 3 . 1 ]
12:    [ 3 2 1 . ]    [ . . 2 ]    [ 3 2 1 . ]
13:    [ 2 3 1 . ]    [ 1 . 2 ]    [ 3 2 . 1 ]
14:    [ 1 3 2 . ]    [ . 1 2 ]    [ 3 . 2 1 ]
15:    [ 1 2 3 . ]    [ 1 1 2 ]    [ 3 . 1 2 ]
16:    [ 2 1 3 . ]    [ . 2 2 ]    [ 3 1 . 2 ]
17:    [ 3 1 2 . ]    [ 1 2 2 ]    [ 3 1 2 . ]
18:    [ 2 . 3 1 ]    [ . . 3 ]    [ 1 3 . 2 ]
19:    [ 3 . 2 1 ]    [ 1 . 3 ]    [ 1 3 2 . ]
20:    [ 3 . 1 2 ]    [ . 1 3 ]    [ 1 2 3 . ]
21:    [ 2 . 1 3 ]    [ 1 1 3 ]    [ 1 2 . 3 ]
22:    [ 1 . 2 3 ]    [ . 2 3 ]    [ 1 . 2 3 ]
23:    [ 1 . 3 2 ]    [ 1 2 3 ]    [ 1 . 3 2 ]
```

Figure 10.10-C: Permutations of 4 elements in single track order starting with the identical permutation.

The ordering obtained is shown in figure 10.10-C, The reversal of the k-the permutation lies at position $(n!)/2 + k$. About 85 million permutations per second can be generated.

10.10.1 Construction of all single track orders

```
112233
231312  <--= permutations of 3 elements in lex order (columns)
323121

000000 112233 231312 323121  <--= concatenate rows and prepend new element

000000 112233 231312 323121  <--= shift  0
323121 000000 112233 231312  <--= shift  6
231312 323121 000000 112233  <--= shift 12
112233 231312 323121 000000  <--= shift 18
```

Figure 10.10-D: Construction of a single track order for permutations of 4 elements from an arbitrary ordering of the permutations of 3 elements.

```
  single track ordering          modified single track ordering
...... 112233 231312 323121       21.113 1.333. .212.1 332.22
323121 ...... 112233 231312       1.333. .212.1 332.22 21.113
231312 323121 ...... 112233       .212.1 332.22 21.113 1.333.
112233 231312 323121 ......       332.22 21.113 1.333. .212.1

^^^^^^                            ^^^^^^
000000                            210321 <--= cyclic shifts
```

Figure 10.10-E: In each of the first $(n-1)!$ permutations in a single track ordering (first block left) an arbitrary rotation can be applied (first block right), leading to a different single track ordering.

A construction for a single track order of $n+1$ elements from an arbitrary ordering of n elements is shown in figure 10.10-D (for $n = 3$ and lexicographic oder). Thereby we obtain as many single track orders for the permutations of n elements as there are orders of the permutations of $n-1$ elements, namely $((n-1)!)!$. One can apply cyclic shifts in each blocks as shown in figure 10.10-E. The shifts in the first $(n-1)!$ positions (first blocks in the figure) determine the shifts for the remaining permutations. Now

there are n different cyclic shifts in each position. Indeed all single track orderings are of this form, so their number is

$$N_s(n) \quad = \quad ((n-1)!)!\; n^{(n-1)!} \tag{10.10-1}$$

The number of single track orders that start with the identical permutation, and where the k-th run of $(n-1)!$ elements starts with k (and thereby all shifts between successive tracks are left shifts by $(n-1)!$ positions) is

$$N_s(n)/n! \quad = \quad ((n-1)!-1)!\; n^{(n-1)!-1} \tag{10.10-2}$$

10.10.2 A single track Gray code

[. 1 2 3 4]	[1 2 3 4 .]	[2 3 4 . 1]	[3 4 . 1 2]	[4 . 1 2 3]
[1 . 2 3 4]	[. 2 3 4 1]	[2 3 4 1 .]	[3 4 1 . 2]	[4 1 . 2 3]
[2 . 1 3 4]	[. 1 3 4 2]	[1 3 4 2 .]	[3 4 2 . 1]	[4 2 . 1 3]
[. 2 1 3 4]	[2 1 3 4 .]	[1 3 4 . 2]	[3 4 . 2 1]	[4 . 2 1 3]
[1 2 . 3 4]	[2 . 3 4 1]	[. 3 4 1 2]	[3 4 1 2 .]	[4 1 2 . 3]
[2 1 . 3 4]	[1 . 3 4 2]	[. 3 4 2 1]	[3 4 2 1 .]	[4 2 1 . 3]
[3 1 . 2 4]	[1 . 2 4 3]	[. 2 4 3 1]	[2 4 3 1 .]	[4 3 1 . 2]
[1 3 . 2 4]	[3 . 2 4 1]	[. 2 4 1 3]	[2 4 1 3 .]	[4 1 3 . 2]
[. 3 1 2 4]	[3 1 2 4 .]	[1 2 4 . 3]	[2 4 . 3 1]	[4 . 3 1 2]
[3 . 1 2 4]	[. 1 2 4 3]	[1 2 4 3 .]	[2 4 3 . 1]	[4 3 . 1 2]
[1 . 3 2 4]	[. 3 2 4 1]	[3 2 4 1 .]	[2 4 1 . 3]	[4 1 . 3 2]
[. 1 3 2 4]	[1 3 2 4 .]	[3 2 4 . 1]	[2 4 . 1 3]	[4 . 1 3 2]
[. 2 3 1 4]	[2 3 1 4 .]	[3 1 4 . 2]	[1 4 . 2 3]	[4 . 2 3 1]
[2 . 3 1 4]	[. 3 1 4 2]	[3 1 4 2 .]	[1 4 2 . 3]	[4 2 . 3 1]
[3 . 2 1 4]	[. 2 1 4 3]	[2 1 4 3 .]	[1 4 3 . 2]	[4 3 . 2 1]
[. 3 2 1 4]	[3 2 1 4 .]	[2 1 4 . 3]	[1 4 . 3 2]	[4 . 3 2 1]
[2 3 . 1 4]	[3 . 1 4 2]	[. 1 4 2 3]	[1 4 2 3 .]	[4 2 3 . 1]
[3 2 . 1 4]	[2 . 1 4 3]	[. 1 4 3 2]	[1 4 3 2 .]	[4 3 2 . 1]
[3 2 1 . 4]	[2 1 . 4 3]	[1 . 4 3 2]	[. 4 3 2 1]	[4 3 2 1 .]
[2 3 1 . 4]	[3 1 . 4 2]	[1 . 4 2 3]	[. 4 2 3 1]	[4 2 3 1 .]
[1 3 2 . 4]	[3 2 . 4 1]	[2 . 4 1 3]	[. 4 1 3 2]	[4 1 3 2 .]
[3 1 2 . 4]	[1 2 . 4 3]	[2 . 4 3 1]	[. 4 3 1 2]	[4 3 1 2 .]
[2 1 3 . 4]	[1 3 . 4 2]	[3 . 4 2 1]	[. 4 2 1 3]	[4 2 1 3 .]
[1 2 3 . 4]	[2 3 . 4 1]	[3 . 4 1 2]	[. 4 1 2 3]	[4 1 2 3 .]

Figure 10.10-F: A cyclic Gray code for the permutations of 5 elements with the single track property.

A Gray code for permutations that has the single track property can be constructed by using a Gray code for the permutations of $n-1$ elements if the first and last permutation are cyclic shifts by one position of each other. Such Gray codes exist for even lengths only. Figure 10.10-F shows a single track Gray code for $n = 5$. For even n we use a Gray code where all but the last element are cyclically shifted between the first and last permutation. Such a Gray code is given in section 10.8.3 on page 263. The resulting single track order is as close to a Gray code as possible, just $n-1$ extra transpositions occur for all permutation of n elements, see figure 10.10-G. The listings were created with the program [FXT: comb/perm-st-gray-demo.cc] which uses [FXT: `class perm_st_gray` in comb/perm-st-gray.h]:

```
1   class perm_st_gray
2   {
3   public:
4       perm_gray_rot1 *G;  // underlying permutations
5
6       ulong *x_;   // permutation
7       ulong *ix_;  // inverse permutation
8       ulong n_;    // number of elements
9       ulong sct_;  // count cyclic shifts
10
11  public:
12      perm_st_gray(ulong n)
13      // Must have n>=2
14      {
15          n_ = (n>=2 ? n : 2);
16          G = new perm_gray_rot1(n-1);
17          x_ = new ulong[n_];
```

```
  1:    [ 0 1 2 3 4 5 ]
  2:    [ 1 0 2 3 4 5 ]
  3:    [ 2 0 1 3 4 5 ]
  4:    [ 0 2 1 3 4 5 ]
  5:    [ 1 2 0 3 4 5 ]
    [--Gray transitions only--]
116:    [ 2 3 1 0 4 5 ]
117:    [ 1 3 2 0 4 5 ]
118:    [ 3 1 2 0 4 5 ]
119:    [ 2 1 3 0 4 5 ]
120:    [ 1 2 3 0 4 5 ]
121:    [ 1 2 3 4 5 0 ]  << (0, 4, 5)
    [--Gray transitions only--]
240:    [ 2 3 0 4 5 1 ]
241:    [ 2 3 4 5 0 1 ]  << (0, 4, 5)
    [--Gray transitions only--]
360:    [ 3 0 4 5 1 2 ]
361:    [ 3 4 5 0 1 2 ]  << (0, 4, 5)
    [--Gray transitions only--]
480:    [ 0 4 5 1 2 3 ]
481:    [ 4 5 0 1 2 3 ]  << (0, 4, 5)
    [--Gray transitions only--]
600:    [ 4 5 1 2 3 0 ]
601:    [ 5 0 1 2 3 4 ]  << (0, 4, 5)
    [--Gray transitions only--]
720:    [ 5 1 2 3 0 4 ]
  1:    [ 0 1 2 3 4 5 ]  << (0, 4, 5)
```

Figure 10.10-G: The single track ordering for odd n with the least number of transpositions contains $n-1$ extra transpositions. Here the non-Gray transitions are cycles between the elements 0, 4, and 5.

```
18          ix_ = new ulong[n_];
19          first();
20      }
21    [--snip--]
22      void first()
23      {
24          G->first();
25          for (ulong j=0; j<n_; ++j)  ix_[j] = x_[j] = j;
26          sct_ = n_;
27      }
```

We define two auxiliary routine to swap elements by their value and by their positions:

```
1   private:
2       void swap_elements(ulong x1, ulong x2)
3       {
4           const ulong i1 = ix_[x1],  i2 = ix_[x2];
5           x_[i1] = x2;   x_[i2] = x1;   // swap2(x_[i1], x_[i2]);
6           ix_[x1] = i2;  ix_[x2] = i1;  // swap2(ix_[x1], ix_[x2]);
7       }
8
9       void swap_positions(ulong i1, ulong i2)
10      {
11          const ulong x1 = x_[i1],  x2 = x_[i2];
12          x_[i1] = x2;   x_[i2] = x1;   // swap2(x_[i1], x_[i2]);
13          ix_[x1] = i2;  ix_[x2] = i1;  // swap2(ix_[x1], ix_[x2]);
14      }
```

The update routine consists of two cases. The frequent case is the update via the underlying permutation:

```
1   public:
2       bool next()
3       {
4           bool q = G->next();
5           if ( q )  // normal update (in underlying permutation of n-1 elements)
6           {
7               ulong i1, i2;  // positions of swaps
8               G->get_swap(i1, i2);
9
10              // rotate positions according to sct:
```

```
11              i1 += sct_;  if ( i1>=n_ )  i1-=n_;
12              i2 += sct_;  if ( i2>=n_ )  i2-=n_;
13
14              swap_positions(i1, i2);
15
16              return true;
17          }
```

The infrequent case happens when the last underlying permutation is encountered:

```
1           else  // goto next cyclic shift (once in (n-1)! updates, n-1 times in total)
2           {
3               G->first();  // restart underlying permutations
4               --sct_;      // adjust cyclic shift
5
6               swap_elements(0, n_-1);
7
8               if ( 0==(n_&1) )  // n even
9                   if ( n_>=4 )  swap_elements(n_-2, n_-1);  // one extra transposition
10
11              return ( 0!=sct_ );
12          }
13      }
```

10.11 Star-transposition order

	permutation	swap	inverse p.
0:	[. 1 2 3]		[. 1 2 3]
1:	[1 . 2 3]	(0, 1)	[1 . 2 3]
2:	[2 . 1 3]	(0, 2)	[1 2 . 3]
3:	[. 2 1 3]	(0, 1)	[. 2 1 3]
4:	[1 2 . 3]	(0, 2)	[2 . 1 3]
5:	[2 1 . 3]	(0, 1)	[2 1 . 3]
6:	[3 1 . 2]	(0, 3)	[2 1 3 .]
7:	[. 1 3 2]	(0, 2)	[. 1 3 2]
8:	[1 . 3 2]	(0, 1)	[1 . 3 2]
9:	[3 . 1 2]	(0, 2)	[1 2 3 .]
10:	[. 3 1 2]	(0, 1)	[. 2 3 1]
11:	[1 3 . 2]	(0, 2)	[2 . 3 1]
12:	[2 3 . 1]	(0, 3)	[2 3 . 1]
13:	[3 2 . 1]	(0, 1)	[2 3 1 .]
14:	[. 2 3 1]	(0, 2)	[. 3 1 2]
15:	[2 . 3 1]	(0, 1)	[1 3 . 2]
16:	[3 . 2 1]	(0, 2)	[1 3 2 .]
17:	[. 3 2 1]	(0, 1)	[. 3 2 1]
18:	[1 3 2 .]	(0, 3)	[3 . 2 1]
19:	[2 3 1 .]	(0, 2)	[3 2 . 1]
20:	[3 2 1 .]	(0, 1)	[3 2 1 .]
21:	[1 2 3 .]	(0, 2)	[3 . 1 2]
22:	[2 1 3 .]	(0, 1)	[3 1 . 2]
23:	[3 1 2 .]	(0, 2)	[3 1 2 .]

Figure 10.11-A: The permutations of 4 elements in star-transposition order. Dots denote zeros.

The permutations can be ordered so that successive permutations differ by a swap of the element at the first position with some other element (star transposition), figure 10.11-A. An algorithm for the generation of such an ordering is given in [170]. The implementation of the algorithm, ascribed to Gideon Ehrlich, is given in [FXT: `class perm_star` in comb/perm-star.h]. The generation of the inverse permutations is an option that is activated by the `#define PERM_STAR_WITH_INVERSE`. If the successive permutations differ by a transposition

```
swap2(a_[0], a_[swp_]);
```

then the inverse can be updated as

```
swap2(ia_[a_[0]], ia_[a_[swp_]]);
```

Note that in the sequence of the inverse permutations the zero is always moved. In the list of the inverse permutations the reversed permutations of the first half are in the second half.

The listing shown in figure 10.11-A can be obtained with [FXT: comb/perm-star-demo.cc]. About 77 million permutations per second are generated and about 115 million when the inverse permutation is not computed. If the only the swaps are of interest then use [FXT: `class perm_star_swaps` in comb/perm-star-swaps.h] whose update routine works at a rate about 158 million per second [FXT: comb/perm-star-swaps-demo.cc].

10.12 Derangement order

```
               permutation        inverse perm.
      0:       [ . 1 2 3 ]        [ . 1 2 3 ]
      1:       [ 3 . 1 2 ]        [ 1 2 3 . ]
      2:       [ 1 2 3 . ]        [ 3 . 1 2 ]
      3:       [ 2 3 . 1 ]        [ 2 3 . 1 ]
      4:       [ 1 . 2 3 ]        [ 1 . 2 3 ]
      5:       [ 3 1 . 2 ]        [ 2 1 3 . ]
      6:       [ . 2 3 1 ]        [ . 3 1 2 ]
      7:       [ 2 3 1 . ]        [ 3 2 . 1 ]
      8:       [ 1 2 . 3 ]        [ 2 . 1 3 ]
      9:       [ 3 1 2 . ]        [ 3 1 2 . ]
     10:       [ 2 . 3 1 ]        [ 1 3 . 2 ]
     11:       [ . 3 1 2 ]        [ . 2 3 1 ]
     12:       [ 2 1 . 3 ]        [ 2 1 . 3 ]
     13:       [ 3 2 1 . ]        [ 3 2 1 . ]
     14:       [ 1 . 3 2 ]        [ 1 . 3 2 ]
     15:       [ . 3 2 1 ]        [ . 3 2 1 ]
     16:       [ 2 . 1 3 ]        [ 1 2 . 3 ]
     17:       [ 3 2 . 1 ]        [ 2 3 1 . ]
     18:       [ . 1 3 2 ]        [ . 1 3 2 ]
     19:       [ 1 3 2 . ]        [ 3 . 2 1 ]
     20:       [ . 2 1 3 ]        [ . 2 1 3 ]
     21:       [ 3 . 2 1 ]        [ 1 3 2 . ]
     22:       [ 2 1 3 . ]        [ 3 1 . 2 ]
     23:       [ 1 3 . 2 ]        [ 2 . 3 1 ]
```

Figure 10.12-A: The permutations of 4 elements in derangement order.

The *derangement order* for permutations is characterized by the fact that two successive permutations have no element at the same position, as shown in figure 10.12-A. The listing was created with the program [FXT: comb/perm-derange-demo.cc]. There is no derangement order for $n = 3$. An implementation of the underlying algorithm (given in [230, p.611]) is [FXT: `class perm_derange` in comb/perm-derange.h]:

```
1   class perm_derange
2   {
3   public:
4       ulong n_;   // number of elements
5       ulong *x_;  // current permutation
6       ulong ctm_; // counter modulo n
7       perm_trotter* T_;
8
9   public:
10      perm_derange(ulong n)
11      // Must have:  n>=4
12      // n=2: trivial, n=3: no solution exists,  n>=4: ok
13      {
14          n_ = n;
15          x_ = new ulong[n_];
16          T_ = new perm_trotter(n_-1);
17          first();
18      }
19    [--snip--]
```

The routine to update the permutation is

```
1       bool next()
```

```
2       {
3           ++ctm_;
4           if ( ctm_>=n_ ) // every n steps: need next perm_trotter
5           {
6               ctm_ = 0;
7               if ( ! T_->next() )  return false;  // current permutation is last
8               const ulong *t = T_->data();
9               for (ulong k=0; k<n_-1; ++k)  x_[k] = t[k];
10              x_[n_-1] = n_-1;  // last element
11          }
12          else  // rotate
13          {
14              if ( ctm_==n_-1 )  rotate_left1(x_, n_);
15              else // last two elements swapped
16              {
17                  rotate_right1(x_, n_);
18                  if ( ctm_==n_-2 )  rotate_right1(x_, n_);
19              }
20          }
21          return true;
22      }
```

The routines `rotate_right1()` and `rotate_last()` rotate the array `x_[]` by one position [FXT: perm/rotate.h]. These rotations are the performance bottleneck, the cost of one update of a length-n permutation is proportional to n. Still, about 35 million permutations per second are generated for $n = 12$.

Gray codes have the minimal number of changes between successive permutations while derangement orders have the maximum. An algorithm for generating all permutations of n objects with k transitions (where $2 \le k \le n$ and $k \ne 3$) is given in [229].

Derangement order for even n *

	permutation	inv. perm.		permutation	inv. perm.	
0:	[. 1 2 3]	[. 1 2 3]	0:	[. 1 2]	[. 1 2]	
1:	[1 2 3 .]	[3 . 1 2]	1:	[1 2 .]	[2 . 1]	
2:	[2 3 . 1]	[2 3 . 1]	2:	[2 . 1]	[1 2 .]	<<
3:	[3 . 1 2]	[1 2 3 .]	3:	[1 . 2]	[1 . 2]	<<
4:	[1 2 . 3]	[2 . 1 3]	4:	[. 2 1]	[. 2 1]	
5:	[2 . 3 1]	[1 3 . 2]	5:	[2 1 .]	[2 1 .]	
6:	[. 3 1 2]	[. 2 3 1]				
7:	[3 1 2 .]	[3 1 2 .]				
8:	[2 . 1 3]	[1 2 . 3]				
9:	[. 1 3 2]	[. 1 3 2]				
10:	[1 3 2 .]	[3 . 2 1]				
11:	[3 2 . 1]	[2 3 1 .]				
12:	[1 . 2 3]	[1 . 2 3]				
13:	[. 2 3 1]	[. 3 1 2]				
14:	[2 3 1 .]	[3 2 . 1]				
15:	[3 1 . 2]	[2 1 3 .]				
16:	[. 2 1 3]	[. 2 1 3]				
17:	[2 1 3 .]	[3 1 . 2]				
18:	[1 3 . 2]	[2 . 3 1]				
19:	[3 . 2 1]	[1 3 2 .]				
20:	[2 1 . 3]	[2 1 . 3]				
21:	[1 . 3 2]	[1 . 3 2]				
22:	[. 3 2 1]	[. 3 2 1]				
23:	[3 2 1 .]	[3 2 1 .]				

Figure 10.12-B: Permutations generated via cyclic shifts. The order is a derangement order for even n (left), but not for odd n (right). Dots denote zeros.

An algorithm for the generation of permutations via cyclic shifts suggested in [177] generates a derangement order if the number n of elements is even, see figure 10.12-B. An implementation of the algorithm, following [170], is [FXT: `class perm_rot` in comb/perm-rot.h]. For odd n the number of times that

```
        ffact    permutation  inv. perm.          ffact      perm.    inv. perm.
 0:  [ . . . ]  [ . 1 2 3 ]  [ . 1 2 3 ]     0:  [ . . ]  [ . 1 2 ]  [ . 1 2 ]
 1:  [ 1 . . ]  [ 3 . 1 2 ]  [ 1 2 3 . ]     1:  [ 1 . ]  [ 2 . 1 ]  [ 1 2 . ]
 2:  [ 2 . . ]  [ 2 3 . 1 ]  [ 2 3 . 1 ]     2:  [ 2 . ]  [ 1 2 . ]  [ 2 . 1 ] <<
 3:  [ 3 . . ]  [ 1 2 3 . ]  [ 3 . 1 2 ]     3:  [ . 1 ]  [ . 2 1 ]  [ . 2 1 ] <<
 4:  [ . 1 . ]  [ . 3 1 2 ]  [ . 2 3 1 ]     4:  [ 1 1 ]  [ 1 . 2 ]  [ 1 . 2 ]
 5:  [ 1 1 . ]  [ 2 . 3 1 ]  [ 1 3 . 2 ]     5:  [ 2 1 ]  [ 2 1 . ]  [ 2 1 . ]
 6:  [ 2 1 . ]  [ 1 2 . 3 ]  [ 2 . 1 3 ]
 7:  [ 3 1 . ]  [ 3 1 2 . ]  [ 3 1 2 . ]
 8:  [ . 2 . ]  [ . 2 3 1 ]  [ . 3 1 2 ]
 9:  [ 1 2 . ]  [ 1 . 2 3 ]  [ 1 . 2 3 ]
10:  [ 2 2 . ]  [ 3 1 . 2 ]  [ 2 1 3 . ]
11:  [ 3 2 . ]  [ 2 3 1 . ]  [ 3 2 . 1 ]
12:  [ . . 1 ]  [ . 1 3 2 ]  [ . 1 3 2 ]
13:  [ 1 . 1 ]  [ 2 . 1 3 ]  [ 1 2 . 3 ]
14:  [ 2 . 1 ]  [ 3 2 . 1 ]  [ 2 3 1 . ]
15:  [ 3 . 1 ]  [ 1 3 2 . ]  [ 3 . 2 1 ]
16:  [ . 1 1 ]  [ . 2 1 3 ]  [ . 2 1 3 ]
17:  [ 1 1 1 ]  [ 3 . 2 1 ]  [ 1 3 2 . ]
18:  [ 2 1 1 ]  [ 1 3 . 2 ]  [ 2 . 3 1 ]
19:  [ 3 1 1 ]  [ 2 1 3 . ]  [ 3 1 . 2 ]
20:  [ . 2 1 ]  [ . 3 2 1 ]  [ . 3 2 1 ]
21:  [ 1 2 1 ]  [ 1 . 3 2 ]  [ 1 . 3 2 ]
22:  [ 2 2 1 ]  [ 2 1 . 3 ]  [ 2 1 . 3 ]
23:  [ 3 2 1 ]  [ 3 2 1 . ]  [ 3 2 1 . ]
```

Figure 10.12-C: Alternative ordering for permutations generated via cyclic shifts. The order is a derangement order for even n (left), but not for odd n (right).

the successor is not a derangement of the predecessor equals $((n+1)/2)! - 1$. The program [FXT: comb/perm-rot-demo.cc] generates the permutations and counts those transitions.

An alternative ordering with the same number of transitions that are no derangements is obtained via mixed radix counting in falling factorial basis and the routine [FXT: comb/perm-rot-unrank-demo.cc]

```
1  void ffact2perm_rot(const ulong *fc, ulong n, ulong *x)
2  // Convert falling factorial number fc[0, ..., n-2] into
3  // permutation of x[0, ... ,n-1].
4  {
5      for (ulong k=0; k<n; ++k)  x[k] = k;
6      for (ulong k=n-1, j=2;  k!=0;  --k, ++j)  rotate_right(x+k-1, j, fc[k-1]);
7  }
```

Figure 10.12-C shows the generated ordering for $n = 4$ and $n = 3$. The observation that the permutations in second ordering are the complemented reversals of the first leads to the unranking routine

```
1  class perm_rot
2  {
3      ulong *a_;  // permutation of n elements
4      ulong n_;
5   [--snip--]
6
7      void goto_ffact(const ulong *d)
8      // Goto permutation corresponding to d[] (i.e. unrank d[]).
9      // d[] must be a valid (falling) factorial mixed radix string.
10     {
11         for (ulong k=0; k<n_; ++k)  a_[k] = k;
12         for (ulong k=n_-1, j=2;  k!=0;  --k, ++j)  rotate_right(a_+k-1, j, d[k-1]);
13         reverse(a_, n_);
14         make_complement(a_, a_, n_);
15     }
16  [--snip--]
17 }
```

Compare to the unranking for permutations by prefix reversals shown in section 10.4.2 on page 248.

	permutation	inverse	ffact-swp
0:	[. 1 2 3]	[. 1 2 3]	[. . .]
1:	[. 1 3 2]	[. 1 3 2]	[. . 1]
2:	[. 2 1 3]	[. 2 1 3]	[. 1 .]
3:	[. 2 3 1]	[. 3 1 2]	[. 1 1]
4:	[. 3 2 1]	[. 3 2 1]	[. 2 .]
5:	[. 3 1 2]	[. 2 3 1]	[. 2 1]
6:	[1 . 2 3]	[1 . 2 3]	[1 . .]
7:	[1 . 3 2]	[1 . 3 2]	[1 . 1]
8:	[1 2 . 3]	[2 . 1 3]	[1 1 .]
9:	[1 2 3 .]	[3 . 1 2]	[1 1 1]
10:	[1 3 2 .]	[3 . 2 1]	[1 2 .]
11:	[1 3 . 2]	[2 . 3 1]	[1 2 1]
12:	[2 1 . 3]	[2 1 . 3]	[2 . .]
13:	[2 1 3 .]	[3 1 . 2]	[2 . 1]
14:	[2 . 1 3]	[1 2 . 3]	[2 1 .]
15:	[2 . 3 1]	[1 3 . 2]	[2 1 1]
16:	[2 3 . 1]	[2 3 . 1]	[2 2 .]
17:	[2 3 1 .]	[3 2 . 1]	[2 2 1]
18:	[3 1 2 .]	[3 1 2 .]	[3 . .]
19:	[3 1 . 2]	[2 1 3 .]	[3 . 1]
20:	[3 2 1 .]	[3 2 1 .]	[3 1 .]
21:	[3 2 . 1]	[2 3 1 .]	[3 1 1]
22:	[3 . 2 1]	[1 3 2 .]	[3 2 .]
23:	[3 . 1 2]	[1 2 3 .]	[3 2 1]

Figure 10.13-A: All permutations of 4 elements (left) and their inverses (middle), and their (swaps-) representations as mixed radix numbers with falling factorial basis. Permutations with common prefixes appear in succession. Dots denote zeros.

10.13 Recursive algorithm for cyclic permutations

A simple recursive algorithm for the generation of the permutations of n elements can be described as follows: Put each of the n element of the array to the first position and generate all permutations of $n-1$ elements. If n equals one, print the permutation.

The order obtained is shown in figure 10.13-A, it corresponds to the alternative (swaps-) factorial representation with falling basis, given in section 10.3.3 on page 241.

The algorithm is implemented in [FXT: `class perm_rec` in comb/perm-rec.h]:

```
1   class perm_rec
2   {
3   public:
4       ulong *x_;   // permutation
5       ulong n_;    // number of elements
6       void (*visit_)(const perm_lex_rec &);  // function to call with each permutation
7
8   public:
9       perm_rec(ulong n)
10      {
11          n_ = n;
12          x_ = new ulong[n_];
13      }
14
15      ~perm_rec()
16      { delete [] x_; }
17
18      void init()
19      {
20          for (ulong k=0; k<n_; ++k)  x_[k] = k;
21      }
22
23      void generate(void (*visit)(const perm_lex_rec &))
24      {
25          visit_ = visit;
26          init();
27          next_rec(0);
28      }
```

The recursive function `next_rec()` is

```
1       void next_rec(ulong d)
2       {
3           if ( d==n_-1 )  visit_(*this);
4           else
5           {
6               const ulong pd = x_[d];
7               for (ulong k=d; k<n_; ++k)
8               {
9
10                  ulong px = x_[k];
11                  x_[k] = pd;  x_[d] = px;  // =^= swap2(x_[d], x_[k]);
12                  next_rec(d+1);
13                  x_[k] = px;  x_[d] = pd;  // =^= swap2(x_[d], x_[k]);
14              }
15          }
16      }
```

The algorithm works because at each recursive call the elements `x[d],...,x[n-1]` are in a different order and when the function returns the elements are in the same order as they were initially. With the '`for`'-statement changed to

```
for (ulong x=n_-1; (long)x>=(long)d; --x)
```

the permutations would appear in reversed order. Changing the loop in the function `next_rec()` to

```
for (ulong k=d; k<n_; ++k)
{
    swap2(x_[d], x_[k]);
    next_rec(d+1, qq);
}
rotate_left1(x_+d, n_-d);
```

produces lexicographic order.

	permutation	cycle	inverse	ffact-swp
0:	[1 2 3 4 .]	(0, 1, 2, 3, 4)	[4 . 1 2 3]	[1 1 1 1]
1:	[1 2 4 . 3]	(0, 1, 2, 4, 3)	[3 . 1 4 2]	[1 1 2 1]
2:	[1 3 . 4 2]	(0, 1, 3, 4, 2)	[2 . 4 1 3]	[1 2 1 1]
3:	[1 3 4 2 .]	(0, 1, 3, 2, 4)	[4 . 3 1 2]	[1 2 2 1]
4:	[1 4 3 . 2]	(0, 1, 4, 2, 3)	[3 . 4 2 1]	[1 3 1 1]
5:	[1 4 . 2 3]	(0, 1, 4, 3, 2)	[2 . 3 4 1]	[1 3 2 1]
6:	[2 . 3 4 1]	(0, 2, 3, 4, 1)	[1 4 . 2 3]	[2 1 1 1]
7:	[2 . 4 1 3]	(0, 2, 4, 3, 1)	[1 3 . 4 2]	[2 1 2 1]
8:	[2 3 1 4 .]	(0, 2, 1, 3, 4)	[4 2 . 1 3]	[2 2 1 1]
9:	[2 3 4 . 1]	(0, 2, 4, 1, 3)	[3 4 . 1 2]	[2 2 2 1]
10:	[2 4 3 1 .]	(0, 2, 3, 1, 4)	[4 3 . 2 1]	[2 3 1 1]
11:	[2 4 1 . 3]	(0, 2, 1, 4, 3)	[3 2 . 4 1]	[2 3 2 1]
12:	[3 2 . 4 1]	(0, 3, 4, 1, 2)	[2 4 1 . 3]	[3 1 1 1]
13:	[3 2 4 1 .]	(0, 3, 1, 2, 4)	[4 3 1 . 2]	[3 1 2 1]
14:	[3 . 1 4 2]	(0, 3, 4, 2, 1)	[1 2 4 . 3]	[3 2 1 1]
15:	[3 . 4 2 1]	(0, 3, 2, 4, 1)	[1 4 3 . 2]	[3 2 2 1]
16:	[3 4 . 1 2]	(0, 3, 1, 4, 2)	[2 3 4 . 1]	[3 3 1 1]
17:	[3 4 1 2 .]	(0, 3, 2, 1, 4)	[4 2 3 . 1]	[3 3 2 1]
18:	[4 2 3 . 1]	(0, 4, 1, 2, 3)	[3 4 1 2 .]	[4 1 1 1]
19:	[4 2 . 1 3]	(0, 4, 3, 1, 2)	[2 3 1 4 .]	[4 1 2 1]
20:	[4 3 1 . 2]	(0, 4, 2, 1, 3)	[3 2 4 1 .]	[4 2 1 1]
21:	[4 3 . 2 1]	(0, 4, 1, 3, 2)	[2 4 3 1 .]	[4 2 2 1]
22:	[4 . 3 1 2]	(0, 4, 2, 3, 1)	[1 3 4 2 .]	[4 3 1 1]
23:	[4 . 1 2 3]	(0, 4, 3, 2, 1)	[1 2 3 4 .]	[4 3 2 1]

Figure 10.13-B: All cyclic permutations of 5 elements and the permutations as cycles, their inverses, and their (swaps-) representations as mixed radix numbers with falling factorial basis (from left to right).

A modified function generates the cyclic permutations (permutations consisting of exactly one cycle of full length, see section 116). The only change is to skip the case $x = d$ in the loop:

```
for (ulong k=d+1; k<n_; ++k)  // omit k==d
```

The cyclic permutations of five elements are shown in figure 10.13-B. The program [FXT: comb/perm-rec-demo.cc] was used to create the figures 10.13-A and 10.13-B.

```
void visit(const perm_rec &P)  // function to call with each permutation
{
    // Print the permutation
}

int
main(int argc, char **argv)
{
    ulong n = 5;  // Number of elements to permute
    bool cq = 1;  // Whether to generate only cyclic permutations
    perm_rec P(n);
    if ( cq )  P.generate_cyclic(visit);
    else       P.generate(visit);
    return 0;
}
```

The routines generate about 57 million permutations and about 37 million cyclic permutations per second.

10.14 Minimal-change order for cyclic permutations

	permutation	fact.num.	cycle
0:	[4 0 1 2 3]	[. . .]	(4, 3, 2, 1, 0,)
1:	[3 4 1 2 0]	[1 . .]	(4, 0, 3, 2, 1,)
2:	[3 0 4 2 1]	[2 . .]	(4, 1, 0, 3, 2,)
3:	[3 0 1 4 2]	[3 . .]	(4, 2, 1, 0, 3,)
4:	[2 3 1 4 0]	[3 1 .]	(4, 0, 2, 1, 3,)
5:	[2 3 4 0 1]	[2 1 .]	(4, 1, 3, 0, 2,)
6:	[2 4 1 0 3]	[1 1 .]	(4, 3, 0, 2, 1,)
7:	[4 3 1 0 2]	[. 1 .]	(4, 2, 1, 3, 0,)
8:	[4 0 3 1 2]	[. 2 .]	(4, 2, 3, 1, 0,)
9:	[2 4 3 1 0]	[1 2 .]	(4, 0, 2, 3, 1,)
10:	[2 0 4 1 3]	[2 2 .]	(4, 3, 1, 0, 2,)
11:	[2 0 3 4 1]	[3 2 .]	(4, 1, 0, 2, 3,)
12:	[1 2 3 4 0]	[3 2 1]	(4, 0, 1, 2, 3,)
13:	[1 2 4 0 3]	[2 2 1]	(4, 3, 0, 1, 2,)
14:	[1 4 3 0 2]	[1 2 1]	(4, 2, 3, 0, 1,)
15:	[4 2 3 0 1]	[. 2 1]	(4, 1, 2, 3, 0,)
16:	[4 3 0 2 1]	[. 1 1]	(4, 1, 3, 2, 0,)
17:	[1 4 0 2 3]	[1 1 1]	(4, 3, 2, 0, 1,)
18:	[1 3 4 2 0]	[2 1 1]	(4, 0, 1, 3, 2,)
19:	[1 3 0 4 2]	[3 1 1]	(4, 2, 0, 1, 3,)
20:	[3 2 0 4 1]	[3 . 1]	(4, 1, 2, 0, 3,)
21:	[3 2 4 1 0]	[2 . 1]	(4, 0, 3, 1, 2,)
22:	[3 4 0 1 2]	[1 . 1]	(4, 2, 0, 3, 1,)
23:	[4 2 0 1 3]	[. . 1]	(4, 3, 1, 2, 0,)

Figure 10.14-A: All cyclic permutations of 5 elements in a minimal-change order.

All cyclic permutations can be generated from a mixed radix Gray code with falling factorial base (see section 9.2 on page 222). Two successive permutations differ at three positions as shown in figure 10.14-A. An constant amortized time (CAT) implementation is [FXT: `class cyclic_perm` in comb/cyclic-perm.h]:

```
1   class cyclic_perm
2   {
3   public:
4       mixedradix_gray *M_;
5       ulong n_;   // number of elements to permute
6       ulong *ix_;  // current permutation (of {0, 1, ..., n-1})
7       ulong *x_;  // inverse permutation
8
9   public:
10      cyclic_perm(ulong n)
```

```
11          : n_(n)
12      {
13          ix_ = new ulong[n_];
14          x_ = new ulong[n_];
15          M_ = new mixedradix_gray(n_-2, 0);  // falling factorial base
16          first();
17      }
18    [--snip--]
```

The computation of the successor uses the position and direction of the mixed radix digit changed with the last increment:

```
1   private:
2       void setup()
3       {
4           const ulong *fc = M_->data();
5           for (ulong k=0; k<n_; ++k)  ix_[k] = k;
6
7           for (ulong k=n_-1; k>1; --k)
8           {
9               ulong z = n_-3-(k-2); // 0, ..., n-3
10              ulong i = fc[z];
11              swap2(ix_[k], ix_[i]);
12          }
13          if ( n_>1 )  swap2(ix_[0], ix_[1]);
14
15          make_inverse(ix_, x_, n_);
16      }
17
18  public:
19      void first()
20      {
21          M_->first();
22          setup();
23      }
24
25      bool next()
26      {
27          if ( false == M_->next() )  { first();  return false; }
28          ulong j = M_->pos();
29
30          if ( j && (x_[0]==n_-1) )  // once in 2*n cases
31          {
32              setup();  // work proportional n
33              // only 3 elements are interchanged
34          }
35          else  // easy case
36          {
37              int d = M_->dir();
38              ulong x2 = (M_->data())[j];
39              ulong x1 = x2 - d,  x3 = n_-1;
40              ulong i1 = x_[x1], i2 = x_[x2], i3 = x_[x3];
41
42              swap2(x_[x1], x_[x2]);
43              swap2(x_[x1], x_[x3]);
44              swap2(ix_[i1], ix_[i2]);
45              swap2(ix_[i2], ix_[i3]);
46          }
47
48          return true;
49      }
50    [--snip--]
```

The order so that the permutation is the same as if one would compute it via the function [FXT: `ffact2cyclic()` in comb/fact2cyclic.cc] which is given in section 10.3.4 on page 243. The listing in figure 10.14-A was created with the program [FXT: comb/cyclic-perm-demo.cc]. About 40 million permutations per second are generated.

10.15 Permutations with special properties

10.15.1 The number of certain permutations

We discuss permutations with special properties, such as involutions, derangements, and permutations with prescribed cycle types.

10.15.1.1 Permutations with m cycles: Stirling cycle numbers

n:	total	m= 1	2	3	4	4	6	7	8	9
1:	1	1								
2:	2	1	1							
3:	6	2	3	1						
4:	24	6	11	6	1					
5:	120	24	50	35	10	1				
6:	720	120	274	225	85	15	1			
7:	5040	720	1764	1624	735	175	21	1		
8:	40320	5040	13068	13132	6769	1960	322	28	1	
9:	362880	40320	109584	118124	67284	22449	4536	546	36	1

Figure 10.15-A: Stirling numbers of the first kind $s(n,m)$ (Stirling cycle numbers).

The number of permutations of n elements into m cycles is given by the (unsigned) *Stirling numbers of the first kind* (or *Stirling cycle numbers*) $s(n,m)$. The first few are shown in figure 10.15-A which was created with the program [FXT: comb/stirling1-demo.cc]. One has $s(1,1)=1$ and

$$s(n,m) \;=\; s(n-1,m-1) + (n-1)\,s(n-1,m) \tag{10.15-1}$$

$$\sum_{m=0}^{n} s(n,m)\,e^m \;=\; \prod_{m=0}^{n} e+m \;=\; e^{\overline{n}} \tag{10.15-2}$$

A generating function is given as relation 35.2-73a on page 695, see also entry A008275 of [240]. The Stirling numbers of the second kind (Stirling set numbers) are treated in section 15.1 on page 336. Many identities involving the Stirling numbers are given in [133, pp.243-253]. We note just a few, writing $S(n,k)$ for the Stirling set numbers:

$$x^n \;=\; \sum_{k=0}^{n} S(n,k)\,x^{\underline{k}} \;=\; \sum_{k=0}^{n} S(n,k)\,(-1)^{n-k}\,x^{\overline{k}} \tag{10.15-3a}$$

where $x^{\underline{k}} = x\,(x-1)\,(x-2)\cdots(x-k+1)$ and $x^{\overline{k}} = x\,(x+1)\,(x+2)\cdots(x+k-1)$.

$$x^{\underline{k}} \;=\; \sum_{k=0}^{n} s(n,k)\,(-1)^{n-k}\,x^k \tag{10.15-3b}$$

$$x^{\overline{k}} \;=\; \sum_{k=0}^{n} s(n,k)\,x^k \tag{10.15-3c}$$

Further [133, p.296], with $D := \frac{d}{dz}$ and $\vartheta = z\frac{d}{dz}$, we have the operator identities

$$\vartheta^n \;=\; \sum_{k=0}^{n} S(n,k)\,z^k\,D^k \tag{10.15-4a}$$

$$z^n\,D^n \;=\; \sum_{k=0}^{n} s(n,k)\,(-1)^{n-k}\,\vartheta^k \tag{10.15-4b}$$

10.15.1.2 Permutations with prescribed cycle type

A permutation of n elements is of *type* $C = [c_1, c_2, c_3, \ldots, c_n]$ if it has c_1 fixed points, c_2 cycles of length 2, c_3 cycles of length 3, and so on. The number $Z_{n,C}$ of permutations of n elements with type C equals [52, p.80]

$$Z_{n,C} \quad = \quad n! \,/\, (c_1!\, c_2!\, c_3! \,\ldots\, c_n!\; 1^{c_1}\, 2^{c_2}\, 3^{c_3} \,\ldots\, n^{c_n}) \tag{10.15-5}$$

We necessarily have $n = 1\,c_1 + 2\,c_2 + \ldots + n\,c_n$, that is, the c_j correspond to a integer partition of n.

The exponential generating function $\exp(L(z))$ where

$$L(z) \quad = \quad \sum_{k=1}^{\infty} \frac{c_k\, z^k}{k} \tag{10.15-6}$$

gives detailed information about the all cycle types:

$$\exp(L(z)) \quad = \quad \sum_{n=0}^{\infty} \left[\sum_{C} \left(Z_{n,C} \prod t_k^{c_k} \right) \right] \frac{z^n}{n!} \tag{10.15-7}$$

That is, the exponent of t_k indicates how many cycles of length k are present in the given cycle type:

```
? n=8;R=O(z^(n+1));
? L=sum(k=1,8,eval(Str("t"k))*z^k/k)+R
  t1*z + 1/2*t2*z^2 + 1/3*t3*z^3 + 1/4*t4*z^4 +  [...]  + 1/8*t8*z^8 + O(z^9)

? serlaplace(exp(L))
 1
 + t1 *z
 + (t1^2 + t2) *z^2
 + (t1^3 + 3*t2*t1 + 2*t3) *z^3
 + (t1^4 + 6*t2*t1^2 + 8*t3*t1 + 3*t2^2 + 6*t4) *z^4
 + (t1^5 + 10*t2*t1^3 + 20*t3*t1^2 + 15*t1*t2^2 + 30*t1*t4 + 20*t3*t2 + 24*t5) *z^5
 + (t1^6 + 15*t2*t1^4 + 40*t3*t1^3 +  [...]  + 15*t2^3 + 90*t4*t2 + 40*t3^2 + 120*t6) *z^6
 + (t1^7 + 21*t2*t1^5 + 70*t3*t1^4 +  [...]  + 504*t5*t2 + 420*t4*t3 + 720*t7) *z^7
 + (t1^8 + 28*t2*t1^6 + 112*t3*t1^5 +  [...]  + 2688*t5*t3 + 1260*t4^2 + 5040*c8) *z^8
 + O(z^9)
```

See section 15.1.2 on page 337 for a similar construction for the EGF for set partitions of given type.

10.15.1.3 Prefix conditions

```
   involutions          up-down          indecomposable       derangements
 1:   1  2  3  4    1:   1  3  2  4    1:   2  3  4  1    1:   2  1  4  3
 2:   1  2  4  3    2:   1  4  2  3    2:   2  4  1  3    2:   2  3  4  1
 3:   1  3  2  4    3:   2  3  1  4    3:   2  4  3  1    3:   2  4  1  3
 4:   1  4  3  2    4:   2  4  1  3    4:   3  1  4  2    4:   3  1  4  2
 5:   2  1  3  4    5:   3  4  1  2    5:   3  2  4  1    5:   3  4  1  2
 6:   2  1  4  3    #perm = 5          6:   3  4  1  2    6:   3  4  2  1
 7:   3  2  1  4                       7:   3  4  2  1    7:   4  1  2  3
 8:   3  4  1  2                       8:   4  1  2  3    8:   4  3  1  2
 9:   4  2  3  1                       9:   4  1  3  2    9:   4  3  2  1
10:   4  3  2  1                      10:   4  2  1  3    #perm = 9
#perm = 10                            11:   4  2  3  1
                                      12:   4  3  1  2
                                      13:   4  3  2  1
                                      #perm = 13
```

Figure 10.15-B: Examples of permutations subject to conditions on the prefixes. From left to right: involutions, up-down permutations, indecomposable permutations and derangements.

Some types of permutations can be generated efficiently by a routine that produces the lexicographically ordered list of permutations subject to conditions for all prefixes. The implementation (following [170]) is [FXT: `class perm_restrpref` in comb/perm-restrpref.h]. The condition (as a function pointer) has to be supplied upon creation of an instance of the class. The program [FXT: comb/perm-restrpref-demo.cc] demonstrates the usage, it can be used to generate all involutions, up-down, indecomposable, or derangement permutations, see figure 10.15-B..

10.15.1.4 Involutions

The sequence $I(n)$ of the number of involutions (self-inverse permutations) starts as ($n \geq 1$)

```
1, 2, 4, 10, 26, 76, 232, 764, 2620, 9496, 35696, 140152, 568504, 2390480, ...
```

This is sequence A000085 of [240]. Compute $I(n)$ using the relation

$$I(n) \quad = \quad I(n-1) + (n-1)\, I(n-2) \tag{10.15-8}$$

```
N=20; v=vector(N);
v[1]=1; v[2]=2;
for(n=3,N,v[n]=v[n-1]+(n-1)*v[n-2]);  v
  [1, 2, 4, 10, 26, 76, ... ]
```

Let $h_n(x)$ be the polynomial such that the coefficient of x^k gives the number of involutions of n elements with k fixed points. The polynomials can be computed recursively via $h_{n+1} = h'_n + x\, h_n$ (starting with $h_0 = 1$). We have $h_n(1) = I(n)$:

```
? h=1;for(k=1,8,h=(deriv(h)+x*h);print(subst(h,x,1),": ",h))
    1: x
    2: x^2 + 1
    4: x^3 + 3*x
   10: x^4 + 6*x^2 + 3
   26: x^5 + 10*x^3 + 15*x
   76: x^6 + 15*x^4 + 45*x^2 + 15
  232: x^7 + 21*x^5 + 105*x^3 + 105*x
  764: x^8 + 28*x^6 + 210*x^4 + 420*x^2 + 105
```

The exponential generating function (EGF) is

$$\sum_{k=0}^{\infty} \frac{I(k)\, x^k}{k!} \quad = \quad \exp\left(x + x^2/2\right) \tag{10.15-9}$$

We further have

$$\sum_{k=0}^{\infty} \frac{h_k(t)\, x^k}{k!} \quad = \quad \exp\left(t\, x + x^2/2\right) \tag{10.15-10}$$

The EGF for the number permutations whose m-th power is identity is given in [275, p.85]:

$$\exp\left(\sum_{d \backslash m} x^d/d\right) \tag{10.15-11}$$

The special case $m = 2$ gives relation 10.15-9.

The condition function for involutions is

```
1   bool cond_inv(const ulong *a, ulong k)
2   {
3       ulong ak = a[k];
4       if ( (ak<=k) && (a[ak]!=k) )  return false;
5       return true;
6   }
```

10.15.1.5 Alternating permutations

The *alternating permutations* (or *up-down permutations*) satisfy $a_0 < a_1 > a_2 < a_3 > \ldots$. The condition function is

```
1   bool cond_updown(const ulong *a, ulong k)
2   // up-down condition: a1 < a2 > a3 < a4 > ...
3   {
4       if ( k<2 ) return true;
5       if ( (k%2) )  return ( a[k]<a[k-1] );
6       else          return ( a[k]>a[k-1] );
7   }
```

Note that the routine is for the permutations of the elements $1, 2, \ldots, n$ in a one-based array.

The sequence $A(n)$ of the number of alternating permutations starts as ($n \geq 1$)

```
1, 1, 2, 5, 16, 61, 272, 1385, 7936, 50521, 353792, 2702765, 22368256, ...
```

It is sequence A000111 of [240], the sequence of the *Euler numbers*. The list can be computed using the relation

$$A(n) \quad = \quad \frac{1}{2} \sum_{k=0}^{n-1} \binom{n-1}{k} A(k)\, A(n-1-k) \tag{10.15-12}$$

```
N=20; v=vector(N+1);
v[0+1]=1;  v[1+1]=1;  v[2+1]=1;  \\ start with zero: v[x] == A(x-1)
for(n=3,N,v[n+1]=1/2*sum(k=0,n-1,binomial(n-1,k)*v[k+1]*v[n-1-k+1]));  v
  [1, 1, 1, 2, 5, 16, 61, 272, ... ]
```

An exponential generating function is

$$\frac{1+\sin(z)}{\cos(z)} \quad = \quad \sum_{k=0}^{\infty} \frac{A(k)\, z^k}{k!} \tag{10.15-13}$$

```
? serlaplace((1+sin(z))/cos(z))
  1 + z + z^2 + 2*z^3 + 5*z^4 + 16*z^5 + 61*z^6 + 272*z^7 + 1385*z^8 + 7936*z^9 + ...
```

10.15.1.6 Indecomposable permutations

The *indecomposable* (or *connected*) permutations satisfy, for $k = 0, 1, \ldots, n-2$, the inequality of sets

$$\{a_0, a_1, \ldots, a_k\} \quad \neq \quad \{0, 1, \ldots, k\} \tag{10.15-14}$$

That is, there is no prefix of length $< n$ that is a permutation of itself. The condition function is

```
1   ulong N;  // set to n in main()
2   bool cond_indecomp(const ulong *a, ulong k)
3   // indecomposable condition: {a1,...,ak} != {1,...,k} for all k<n
4   {
5       if ( k==N )  return true;
6       for (ulong i=1; i<=k; ++i)  if ( a[i]>k )  return true;
7       return false;
8   }
```

The sequence $C(n)$ of the number of indecomposable permutations starts as ($n \geq 1$)

```
1, 1, 3, 13, 71, 461, 3447, 29093, 273343, 2829325, 31998903, 392743957, ...
```

This is sequence A003319 of [240]. Compute $C(n)$ using

$$C(n) \quad = \quad n! - \sum_{k=1}^{n-1} k!\, C(n-k) \tag{10.15-15}$$

```
N=20; v=vector(N);
for(n=1,N,v[n]=n!-sum(k=1,n-1,k!*v[n-k]));  v
  [1, 1, 3, 13, 71, 461, 3447, ... ]
```

The ordinary generating function can be given as

$$\sum_{n=1}^{\infty} C(n)\, z^n \quad = \quad 1 - \frac{1}{\sum_{k=0}^{\infty} k!\, z^k} \;=\; z + z^2 + 3\, z^3 + 13\, z^4 + 71\, z^5 + \ldots \tag{10.15-16}$$

The logarithmic generating function is

$$\sum_{n=1}^{\infty} \frac{C(n)}{n}\, z^n \quad = \quad z \left[1 + \log\left(\sum_{k=0}^{\infty} k!\, z^k \right) \right] \;=\; z + z^2 + \frac{3}{2}\, z^3 + \frac{13}{3}\, z^4 + \frac{71}{4}\, z^5 + \ldots \tag{10.15-17}$$

10.15.1.7 Derangements

A permutation is a *derangement* if $a_k \neq k$ for all k:

```
1   bool cond_derange(const ulong *a, ulong k)
2   // derangement condition: f[k]!=k for all k
3   {
4       return ( a[k]!=k );
5   }
```

The sequence $D(n)$ of the number of derangements starts as ($n \geq 1$)

```
0, 1, 2, 9, 44, 265, 1854, 14833, 133496, 1334961, 14684570, 176214841, ...
```

This is sequence A000166 of [240], the *subfactorial numbers*. Compute $D(n)$ using either of

$$\begin{aligned} D(n) &= (n-1)\,[D(n-1)+D(n-2)] && \text{(10.15-18a)}\\ &= n\,D(n-1)+(-1)^n && \text{(10.15-18b)}\\ &= \sum_{k=0}^{n}(-1)^{n-k}\,\frac{n!}{(n-k)!} && \text{(10.15-18c)}\\ &= \lfloor (n!+1/2)/e \rfloor && \text{(10.15-18d)} \end{aligned}$$

where $e = \exp(1)$. We use the recursion 10.15-18a:

```
N=20;  v=vector(N);  v[1]=0;  v[2]=1;
for(n=3,N,v[n]=(n-1)*(v[n-1]+v[n-2]));  v
  [0, 1, 2, 9, 44, 265, 1854, 14833, ... ]
```

The exponential generating function can be obtained by setting $t_1 = 0$ and $t_k = 1$ for $k \neq 1$ in relation 10.15-6: we have $L(z) = \log 1/(1-z) - x$ and

$$\sum_{k=0}^{\infty}\frac{D(n)\,z^n}{n!} \;=\; \exp L(z) \;=\; \frac{\exp(-z)}{1-z} \qquad \text{(10.15-19)}$$

10.15.2 Permutations with distance restrictions

We present constructions for Gray codes for permutations with certain restrictions. These are computed from Gray codes of mixed radix numbers with factorial basis. We write $p(k)$ for the position of the element k in a given permutation.

10.15.2.1 Permutations where $p(k) \leq k+1$

Let $M(n)$ the number of permutations of n elements where no element can move more than one place to the right. We have $M(n) = 2^{n-1}$. A Gray code for these permutation is shown in figure 10.15-C which was created with the program [FXT: comb/perm-right1-gray-demo.cc]. $M(n)$ also counts the permutations that start as a rising sequence (ending in the maximal element) and end as a falling sequence. The recursion for the array in the leftmost column of figure 10.15-C can be obtained by the recursion

```
1   void Y_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           if ( z )  // forward:
7           {
8               // words 0, 10, 200, 3000, 40000, ...
9               ulong k = 0;
10              do
11              {
12                  ff[d] = k;
13                  Y_rec(d+k+1, !z);
14              }
```

```
             ffact        perm            inv. perm       ffact(inv)
      1:     . 3 . .    [ 0 4 1 2 3 ]    [ 0 2 3 4 1 ]     . 1 1 1
      2:     . 2 . .    [ 0 3 1 2 4 ]    [ 0 2 3 1 4 ]     . 1 1 .
      3:     . 1 . .    [ 0 2 1 3 4 ]    [ 0 2 1 3 4 ]     . 1 . .
      4:     . 1 . 1    [ 0 2 1 4 3 ]    [ 0 2 1 4 3 ]     . 1 . 1
      5:     . . . 1    [ 0 1 2 4 3 ]    [ 0 1 2 4 3 ]     . . . 1
      6:     . . . .    [ 0 1 2 3 4 ]    [ 0 1 2 3 4 ]     . . . .
      7:     . . 1 .    [ 0 1 3 2 4 ]    [ 0 1 3 2 4 ]     . . 1 .
      8:     . . 2 .    [ 0 1 4 2 3 ]    [ 0 1 3 4 2 ]     . . 1 1
      9:     1 . 2 .    [ 1 0 4 2 3 ]    [ 1 0 3 4 2 ]     1 . 1 1
     10:     1 . 1 .    [ 1 0 3 2 4 ]    [ 1 0 3 2 4 ]     1 . 1 .
     11:     1 . . .    [ 1 0 2 3 4 ]    [ 1 0 2 3 4 ]     1 . . .
     12:     1 . . 1    [ 1 0 2 4 3 ]    [ 1 0 2 4 3 ]     1 . . 1
     13:     2 . . 1    [ 2 0 1 4 3 ]    [ 1 2 0 4 3 ]     1 1 . 1
     14:     2 . . .    [ 2 0 1 3 4 ]    [ 1 2 0 3 4 ]     1 1 . .
     15:     3 . . .    [ 3 0 1 2 4 ]    [ 1 2 3 0 4 ]     1 1 1 .
     16:     4 . . .    [ 4 0 1 2 3 ]    [ 1 2 3 4 0 ]     1 1 1 1
```

Figure 10.15-C: Gray code for the permutations of 5 elements where no element lies more than one place to the right of its position in the identical permutation.

```
15              while ( ++k <= (n-d) );
16          }
17          else      // backward:
18          {
19              // words ..., 40000, 3000, 200, 10, 0
20              ulong k = n-d+1;
21              do
22              {
23                  --k;
24                  ff[d] = k;
25                  Y_rec(d+k+1, !z);
26              }
27              while ( k != 0 );
28          }
29      }
30  }
```

The array `ff` (of length `n`) must be initialized with zeros and the initial call is `Y_rec(0, true);`. About 85 million words per second are generated. In the inverse permutations (where no element is more than one place left of its original position) the swaps are adjacent and their position is determined by the ruler function. Thereby the inverse permutations can be generated using [FXT: `class ruler_func` in comb/ruler-func.h] which described in section 8.2.3 on page 206.

10.15.2.2 Permutations where $k - 1 \le p(k) \le k + 1$

```
          ffact          perm                      ffact          perm
     1:  1 . . 1 . .  [ 1 0 2 4 3 5 6 ]      14:  . . . . . 1  [ 0 1 2 3 4 6 5 ]
     2:  1 . . 1 . 1  [ 1 0 2 4 3 6 5 ]      15:  . . . 1 . 1  [ 0 1 2 4 3 6 5 ]
     3:  1 . . . . 1  [ 1 0 2 3 4 6 5 ]      16:  . . . 1 . .  [ 0 1 2 4 3 5 6 ]
     4:  1 . . . . .  [ 1 0 2 3 4 5 6 ]      17:  . 1 . 1 . .  [ 0 2 1 4 3 5 6 ]
     5:  1 . . . 1 .  [ 1 0 2 3 5 4 6 ]      18:  . 1 . 1 . 1  [ 0 2 1 4 3 6 5 ]
     6:  1 . 1 . 1 .  [ 1 0 3 2 5 4 6 ]      19:  . 1 . . . 1  [ 0 2 1 3 4 6 5 ]
     7:  1 . 1 . . .  [ 1 0 3 2 4 5 6 ]      20:  . 1 . . . .  [ 0 2 1 3 4 5 6 ]
     8:  1 . 1 . . 1  [ 1 0 3 2 4 6 5 ]      21:  . 1 . . 1 .  [ 0 2 1 3 5 4 6 ]
     9:  . . 1 . . 1  [ 0 1 3 2 4 6 5 ]
    10:  . . 1 . . .  [ 0 1 3 2 4 5 6 ]
    11:  . . 1 . 1 .  [ 0 1 3 2 5 4 6 ]
    12:  . . . . 1 .  [ 0 1 2 3 5 4 6 ]
    13:  . . . . . .  [ 0 1 2 3 4 5 6 ]
```

Figure 10.15-D: Gray code for the permutations of 7 elements where no element lies more than one place away from its position in the identical permutation. The permutations are self-inverse.

Let $F(n)$ the number of permutations of n elements where no element can move more than one place to

the left. Then $F(n)$ is the $(n+1)$-st Fibonacci number. A Gray code for these permutation is shown in figure 10.15-D which was created with the program [FXT: comb/perm-dist1-gray-demo.cc].

10.15.2.3 Permutations where $k-1 \le p(k) \le k+d$

```
            ffact          perm              inv. perm        ffact(inv)
 1:     1 1 . . 1     [ 1 2 0 3 5 4 ]     [ 2 0 1 3 5 4 ]     2 . . . 1
 2:     1 1 . . .     [ 1 2 0 3 4 5 ]     [ 2 0 1 3 4 5 ]     2 . . . .
 3:     1 1 . 1 .     [ 1 2 0 4 3 5 ]     [ 2 0 1 4 3 5 ]     2 . . 1 .
 4:     1 1 . 1 1     [ 1 2 0 4 5 3 ]     [ 2 0 1 5 3 4 ]     2 . . 2 .
 5:     1 . . 1 1     [ 1 0 2 4 5 3 ]     [ 1 0 2 5 3 4 ]     1 . . 2 .
 6:     1 . . 1 .     [ 1 0 2 4 3 5 ]     [ 1 0 2 4 3 5 ]     1 . . 1 .
 7:     1 . . . .     [ 1 0 2 3 4 5 ]     [ 1 0 2 3 4 5 ]     1 . . . .
 8:     1 . . . 1     [ 1 0 2 3 5 4 ]     [ 1 0 2 3 5 4 ]     1 . . . 1
 9:     1 . 1 . 1     [ 1 0 3 2 5 4 ]     [ 1 0 3 2 5 4 ]     1 . 1 . 1
10:     1 . 1 . .     [ 1 0 3 2 4 5 ]     [ 1 0 3 2 4 5 ]     1 . 1 . .
11:     1 . 1 1 .     [ 1 0 3 4 2 5 ]     [ 1 0 4 2 3 5 ]     1 . 2 . .
12:     . . 1 1 .     [ 0 1 3 4 2 5 ]     [ 0 1 4 2 3 5 ]     . . 2 . .
13:     . . 1 . .     [ 0 1 3 2 4 5 ]     [ 0 1 3 2 4 5 ]     . . 1 . .
14:     . . 1 . 1     [ 0 1 3 2 5 4 ]     [ 0 1 3 2 5 4 ]     . . 1 . 1
15:     . . . . 1     [ 0 1 2 3 5 4 ]     [ 0 1 2 3 5 4 ]     . . . . 1
16:     . . . . .     [ 0 1 2 3 4 5 ]     [ 0 1 2 3 4 5 ]     . . . . .
17:     . . . 1 .     [ 0 1 2 4 3 5 ]     [ 0 1 2 4 3 5 ]     . . . 1 .
18:     . . . 1 1     [ 0 1 2 4 5 3 ]     [ 0 1 2 5 3 4 ]     . . . 2 .
19:     . 1 . 1 1     [ 0 2 1 4 5 3 ]     [ 0 2 1 5 3 4 ]     . 1 . 2 .
20:     . 1 . 1 .     [ 0 2 1 4 3 5 ]     [ 0 2 1 4 3 5 ]     . 1 . 1 .
21:     . 1 . . .     [ 0 2 1 3 4 5 ]     [ 0 2 1 3 4 5 ]     . 1 . . .
22:     . 1 . . 1     [ 0 2 1 3 5 4 ]     [ 0 2 1 3 5 4 ]     . 1 . . 1
23:     . 1 1 . 1     [ 0 2 3 1 5 4 ]     [ 0 3 1 2 5 4 ]     . 2 . . 1
24:     . 1 1 . .     [ 0 2 3 1 4 5 ]     [ 0 3 1 2 4 5 ]     . 2 . . .
```

Figure 10.15-E: Gray code for the permutations of 6 elements where no element lies more than one place to the left or two places to the right of its position in the identical permutation.

A Gray code for the permutations where no element lies more than one place to the left or d places to the right of its original position can be obtained via the Gray codes for binary words with at most d successive ones given in section 12.2 on page 298. Figure 10.15-E shows the permutations of 6 elements with $d = 2$, it was created with the program [FXT: comb/perm-l1r2-gray-demo.cc]. The array shown leftmost in figure 10.15-E can be generated via the recursion

```
1   void Y_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           const ulong w = n-d;
7           if ( z )
8           {
9               if ( w>1 )  { ff[d]=1; ff[d+1]=1; ff[d+2]=0;  Y_rec(d+3, !z); }
10              ff[d]=1; ff[d+1]=0;  Y_rec(d+2, !z);
11              ff[d]=0;  Y_rec(d+1, !z);
12          }
13          else
14          {
15              ff[d]=0;  Y_rec(d+1, !z);
16              ff[d]=1; ff[d+1]=0;  Y_rec(d+2, !z);
17              if ( w>1 )  { ff[d]=1; ff[d+1]=1; ff[d+2]=0; Y_rec(d+3, !z); }
18          }
19      }
20  }
```

If the two lines starting `if ( w>1 )` are omitted then the Fibonacci words are obtained. About 110 million words per second are generated.

Chapter 11

Subsets and permutations of a multiset

11.1 Subsets of a multiset

```
n == 630
primes    = [   2   3   5   7 ]
exponents = [   1   2   1   1 ]

           d      auxiliary products          exponents    change @
 1:        1     [   1   1   1   1   1 ]     [ . . . . ]       4
 2:        2     [   2   1   1   1   1 ]     [ 1 . . . ]       0
 3:        3     [   3   3   1   1   1 ]     [ . 1 . . ]       1
 4:        6     [   6   3   1   1   1 ]     [ 1 1 . . ]       0
 5:        9     [   9   9   1   1   1 ]     [ . 2 . . ]       1
 6:       18     [  18   9   1   1   1 ]     [ 1 2 . . ]       0
 7:        5     [   5   5   5   1   1 ]     [ . . 1 . ]       2
 8:       10     [  10   5   5   1   1 ]     [ 1 . 1 . ]       0
 9:       15     [  15  15   5   1   1 ]     [ . 1 1 . ]       1
10:       30     [  30  15   5   1   1 ]     [ 1 1 1 . ]       0
11:       45     [  45  45   5   1   1 ]     [ . 2 1 . ]       1
12:       90     [  90  45   5   1   1 ]     [ 1 2 1 . ]       0
13:        7     [   7   7   7   7   1 ]     [ . . . 1 ]       3
14:       14     [  14   7   7   7   1 ]     [ 1 . . 1 ]       0
15:       21     [  21  21   7   7   1 ]     [ . 1 . 1 ]       1
16:       42     [  42  21   7   7   1 ]     [ 1 1 . 1 ]       0
17:       63     [  63  63   7   7   1 ]     [ . 2 . 1 ]       1
18:      126     [ 126  63   7   7   1 ]     [ 1 2 . 1 ]       0
19:       35     [  35  35  35   7   1 ]     [ . . 1 1 ]       2
20:       70     [  70  35  35   7   1 ]     [ 1 . 1 1 ]       0
21:      105     [ 105 105  35   7   1 ]     [ . 1 1 1 ]       1
22:      210     [ 210 105  35   7   1 ]     [ 1 1 1 1 ]       0
23:      315     [ 315 315  35   7   1 ]     [ . 2 1 1 ]       1
24:      630     [ 630 315  35   7   1 ]     [ 1 2 1 1 ]       0
```

Figure 11.1-A: Divisors of $630 = 2^1 \cdot 3^2 \cdot 5^1 \cdot 7^1$ generated as subsets of the multiset of exponents.

A *multiset* is set where elements can be repeated. A subset of a set of n elements can be identified with the bits of all n-bit binary words. The subsets of a multiset can be obtained as mixed radix numbers: if the j-th element is repeated r_j times then the radix of digit j has to be $r_j + 1$. Thereby all methods of chapter 9 on page 219 can be applied.

As an example, all divisors of a number x whose factorization $x = p_0^{e_0} \cdot p_1^{e_1} \cdots p_{n-1}^{e_{n-1}}$ is known can be obtained via the length-n mixed radix numbers with radices $[e_0 + 1, e_1 + 1, \ldots, e_{n-1} + 1]$. The implementation [FXT: `class divisors` in mod/divisors.h] generates the subsets of the exponent-multiset

in counting order (figure 11.1-A shows the data for $x = 630$). An auxiliary array T of products is updated with each step: if the changed digit (at position j) became 1 then set $t := T_{j+1} \cdot p_j$, else set $t := T_j \cdot p_j$. Set $T_i = t$ for all $0 \le i \le j$. A sentinel element $T_n = 1$ avoids unnecessary code. Figure 11.1-A was created with the program [FXT: mod/divisors-demo.cc]. The computation of all products of k out of n given factors is described in section 6.1.2 on page 178.

11.2 Permutations of a multiset

```
        (2, 2, 1)                (6, 2)                (1, 1, 1, 1)
 1:  [ . . 1 1 2 ]     1:  [ . . . . . . 1 1 ]     1:  [ . 1 2 3 ]
 2:  [ . . 1 2 1 ]     2:  [ . . . . . 1 . 1 ]     2:  [ . 1 3 2 ]
 3:  [ . . 2 1 1 ]     3:  [ . . . . . 1 1 . ]     3:  [ . 2 1 3 ]
 4:  [ . 1 . 1 2 ]     4:  [ . . . . 1 . . 1 ]     4:  [ . 2 3 1 ]
 5:  [ . 1 . 2 1 ]     5:  [ . . . . 1 . 1 . ]     5:  [ . 3 1 2 ]
 6:  [ . 1 1 . 2 ]     6:  [ . . . . 1 1 . . ]     6:  [ . 3 2 1 ]
 7:  [ . 1 1 2 . ]     7:  [ . . . 1 . . . 1 ]     7:  [ 1 . 2 3 ]
 8:  [ . 1 2 . 1 ]     8:  [ . . . 1 . . 1 . ]     8:  [ 1 . 3 2 ]
 9:  [ . 1 2 1 . ]     9:  [ . . . 1 . 1 . . ]     9:  [ 1 2 . 3 ]
10:  [ . 2 . 1 1 ]    10:  [ . . . 1 1 . . . ]    10:  [ 1 2 3 . ]
11:  [ . 2 1 . 1 ]    11:  [ . . 1 . . . . 1 ]    11:  [ 1 3 . 2 ]
12:  [ . 2 1 1 . ]    12:  [ . . 1 . . . 1 . ]    12:  [ 1 3 2 . ]
13:  [ 1 . . 1 2 ]    13:  [ . . 1 . . 1 . . ]    13:  [ 2 . 1 3 ]
14:  [ 1 . . 2 1 ]    14:  [ . . 1 . 1 . . . ]    14:  [ 2 . 3 1 ]
15:  [ 1 . 1 . 2 ]    15:  [ . . 1 1 . . . . ]    15:  [ 2 1 . 3 ]
16:  [ 1 . 1 2 . ]    16:  [ . 1 . . . . . 1 ]    16:  [ 2 1 3 . ]
17:  [ 1 . 2 . 1 ]    17:  [ . 1 . . . . 1 . ]    17:  [ 2 3 . 1 ]
18:  [ 1 . 2 1 . ]    18:  [ . 1 . . . 1 . . ]    18:  [ 2 3 1 . ]
19:  [ 1 1 . . 2 ]    19:  [ . 1 . . 1 . . . ]    19:  [ 3 . 1 2 ]
20:  [ 1 1 . 2 . ]    20:  [ . 1 . 1 . . . . ]    20:  [ 3 . 2 1 ]
21:  [ 1 1 2 . . ]    21:  [ . 1 1 . . . . . ]    21:  [ 3 1 . 2 ]
22:  [ 1 2 . . 1 ]    22:  [ 1 . . . . . . 1 ]    22:  [ 3 1 2 . ]
23:  [ 1 2 . 1 . ]    23:  [ 1 . . . . . 1 . ]    23:  [ 3 2 . 1 ]
24:  [ 1 2 1 . . ]    24:  [ 1 . . . . 1 . . ]    24:  [ 3 2 1 . ]
25:  [ 2 . . 1 1 ]    25:  [ 1 . . . 1 . . . ]
26:  [ 2 . 1 . 1 ]    26:  [ 1 . . 1 . . . . ]
27:  [ 2 . 1 1 . ]    27:  [ 1 . 1 . . . . . ]
28:  [ 2 1 . . 1 ]    28:  [ 1 1 . . . . . . ]
29:  [ 2 1 . 1 . ]
30:  [ 2 1 1 . . ]
```

Figure 11.2-A: Permutations of multisets in lexicographic order: the multiset $(2, 2, 1)$ (left), $(6, 2)$ (combinations $\binom{6+2}{2}$), middle), and $(1, 1, 1, 1)$ (permutations of four elements, right). Dots denote zeros.

We write $(r_0, r_1, \ldots, r_{k-1})$ for a multiset with r_0 elements of the first sort, r_1 of the second sort, ..., r_{k-1} elements of the k-th sort. The total number of elements is $n = \sum_{j=0}^{k-1} r_k$. For the elements of the j-th sort we always use the number j. The number of permutations $P(r_0, r_1, \ldots, r_{k-1})$ of the multiset $(r_0, r_1, \ldots, r_{k-1})$ is a *multinomial coefficient*:

$$P(r_0, r_1, \ldots, r_{k-1}) \;=\; \binom{n}{r_0,\, r_1,\, r_2, \ldots,\, r_{k-1}} \;=\; \frac{n!}{r_0! \cdots r_{k-1}!} \tag{11.2-1a}$$

$$= \binom{n}{r_0} \binom{n-r_0}{r_1} \binom{n-r_0-r_1}{r_2} \cdots \binom{r_{k-3}+r_{k-2}+r_{k-1}}{r_{k-3}} \binom{r_{k-2}+r_{k-1}}{r_{k-2}} \tag{11.2-1b}$$

Relation 11.2-1a is obtained by observing that among the $n!$ ways to arrange all n elements $r_0!$ permutations of the first sort of elements, $r_1!$ of the second, and so on, lead to identical permutations.

11.2.1 Recursive generation

Let $[r_0, r_1, r_2, \ldots, r_{k-1}]$ denote the list of all permutations of the multiset $(r_0, r_1, r_2, \ldots, r_{k-1})$. The recursion

$$[r_0, r_1, r_2, \ldots, r_{k-1}] \quad = \quad \begin{matrix} r_0 \,.\, [r_0 - 1, r_1, r_2, \ldots, r_{k-1}] \\ r_1 \,.\, [r_0, r_1 - 1, r_2, \ldots, r_{k-1}] \\ r_2 \,.\, [r_0, r_1, r_2 - 1, \ldots, r_{k-1}] \\ \vdots \\ r_{k-1} \,.\, [r_0, r_1, r_2, \ldots, r_{k-1} - 1] \end{matrix} \tag{11.2-2}$$

is used to obtain the following procedure [FXT: comb/mset-lex-rec-demo.cc]:

```
ulong n;    // number of objects
ulong *ms;  // multiset data in ms[0], ..., ms[n-1]
ulong k;    // number of different sorts of objects
ulong *r;   // number of elements '0' in r[0], '1' in r[1], ..., 'k-1' in r[k-1]
```

With the recursion

```
void mset_rec(ulong d)
{
    if ( d>=n )  visit();
    else
    {
        for (ulong j=0; j<k; ++j)  // for all buckets
        {
            ++wct;
            if ( r[j] )    // bucket has elements left
            {
                ++rct;
                --r[j];     // take element from bucket
                ms[d] = j;  // put element in place
                mset_rec(d+1);  // recursion
                ++r[j];     // put element back
            }
        }
    }
}
```

and the initial call `mset_rec(0)` we generate all multiset permutations in lexicographic order. As given the routine is inefficient when used with (many) small numbers r_j. An extreme case is $r_j = 1$ for all j, corresponding to the (regular) permutations: we have $n = k$ and for the $n!$ permutations the work is proportional to n^n. The method can be made efficient by maintaining a list of pointer to the next nonzero 'bucket' `nk[]` [FXT: `class mset_lex_rec` in comb/mset-lex-rec.h]:

```
class mset_lex_rec
{
public:
    ulong k_;    // number of different sorts of objects
    ulong *r_;   // number of elements '0' in r[0], '1' in r[1], ..., 'k-1' in r[k-1]
    ulong n_;    // number of objects
    ulong *ms_;  // multiset data in ms[0], ..., ms[n-1]
    ulong *nn_;  // position of next nonempty bucket
    void (*visit_)(const mset_lex_rec &);  // function to call with each permutation
    ulong ct_;  // count objects
    ulong rct_; // count recursions (==work)
  [--snip--]
```

The initializer takes as arguments an array of multiplicities and its length:

```
public:
    mset_lex_rec(ulong *r, ulong k)
    {
        k_ = k;
        r_ = new ulong[k];
        for (ulong j=0; j<k_; ++j)  r_[j] = r[j];  // get buckets

        n_ = 0;
        for (ulong j=0; j<k_; ++j)  n_ += r_[j];
        ms_ = new ulong[n_];
```

```
11
12          nn_ = new ulong[k_+1];  // incl sentinel
13          for (ulong j=0; j<k_; ++j)  nn_[j] = j+1;
14          nn_[k] = 0;  // pointer to first nonempty bucket
15      }
16    [--snip--]
```

The method to generate all permutations takes a 'visit' function as argument:

```
1       void generate(void (*visit)(const mset_lex_rec &))
2       {
3           visit_ = visit;
4           ct_ = 0;
5           rct_ = 0;
6           mset_rec(0);
7       }
8
9   private:
10      void mset_rec(ulong d);
11  };
```

The recursion itself is [FXT: comb/mset-lex-rec.cc]:

```
1   void mset_lex_rec::mset_rec(ulong d)
2   {
3       if ( d>=n_ )
4       {
5           ++ct_;
6           visit_( *this );
7       }
8       else
9       {
10          for (ulong jf=k_, j=nn_[jf];  j<k_;  jf=j, j=nn_[j])  // for all nonempty buckets
11          {
12              ++rct_;  // work == number of recursions
13
14              --r_[j];     // take element from bucket
15              ms_[d] = j;  // put element in place
16
17              if ( r_[j]==0 )  // bucket now empty?
18              {
19                  ulong f = nn_[jf]; // where we come from
20                  nn_[jf] = nn_[j];  // let recursions skip over j
21                  mset_rec(d+1);     // recursion
22                  nn_[jf] = f;       // remove skip
23              }
24              else  mset_rec(d+1);   // recursion
25
26              ++r_[j];      // put element back
27          }
28      }
29  }
```

Note that the test whether the current bucket is nonempty is omitted as empty buckets are skipped. Now the work involved with (regular) permutations is (less than) $e = 2.71828\ldots$ times the number of the generated permutations. Usage of the class is shown in [FXT: comb/mset-lex-rec2-demo.cc]. The permutations of 12 elements are generated at a rate of about 25 million per second, the combinations $\binom{30}{15}$ at about 40 million per second, and the permutations of $(2, 2, 2, 3, 3, 3)$ at about 20 million per second.

11.2.2 Iterative generation

The algorithm to generate the next permutation in lexicographic order given in section 10.1 on page 234 can be adapted for an iterative method for multiset permutations [FXT: `class mset_lex` in comb/mset-lex.h]:

```
1   class mset_lex
2   {
3   public:
4       ulong k_;    // number of different sorts of objects
5       ulong *r_;   // number of elements '0' in r[0], '1' in r[1], ..., 'k-1' in r[k-1]
6       ulong n_;    // number of objects
7       ulong *ms_;  // multiset data in ms[0], ..., ms[n-1], sentinel at [-1]
```

```
 8
 9  public:
10      mset_lex(const ulong *r, ulong k)
11      {
12          k_ = k;
13          r_ = new ulong[k];
14          for (ulong j=0; j<k_; ++j)  r_[j] = r[j];  // get buckets
15
16          n_ = 0;
17          for (ulong j=0; j<k_; ++j)  n_ += r_[j];
18          ms_ = new ulong[n_+1];
19          ms_[0] = 0;  // sentinel
20          ++ms_;  // nota bene
21
22          first();
23      }
24
25      void first()
26      {
27          for (ulong j=0, i=0;  j<k_;  ++j)
28              for (ulong h=r_[j];  h!=0;  --h, ++i)  ms_[i] = j;
29      }
30    [--snip--]
```

The only change in the update routine is to replace the operators > by >= in the scanning loops:

```
1       bool next()
2       {
3           // find for rightmost pair with p_[i] < p_[i+1]:
4           const ulong n1 = n_ - 1;
5           ulong i = n1;
6           do  { --i; }  while ( ms_[i] >= ms_[i+1] );  // can touch sentinel
7           if ( (long)i<0 )  return false;  // last sequence is falling seq.
8
9           // find rightmost element p[j] smaller than p[i]:
10          ulong j = n1;
11          while ( ms_[i] >= ms_[j] )  { --j; }
12
13          swap2(ms_[i], ms_[j]);
14
15          // Here the elements p[i+1], ..., p[n-1] are a falling sequence.
16          // Reverse order to the right:
17          ulong r = n1;
18          ulong s = i + 1;
19          while ( r > s )  { swap2(ms_[r], ms_[s]);  --r;  ++s; }
20
21          return true;
22      }
23  }
```

Usage of the class is shown in [FXT: comb/mset-lex-demo.cc]:

```
ulong ct = 0;
do
{
    // visit
}
while ( P.next() );
```

The permutations of 12 elements are generated at a rate of about 110 million per second, the combinations $\binom{30}{15}$ at about 60 million per second, and the permutations of $(2, 2, 2, 3, 3, 3)$ at about 82 million per second.

Chapter 12

Gray codes for strings with restrictions

We give constructions for Gray codes for strings with certain restrictions, such as forbidding two successive zeros or nonzero digits. The constraints considered are such that the number of strings of a given type satisfies a linear recursion with constant coefficients.

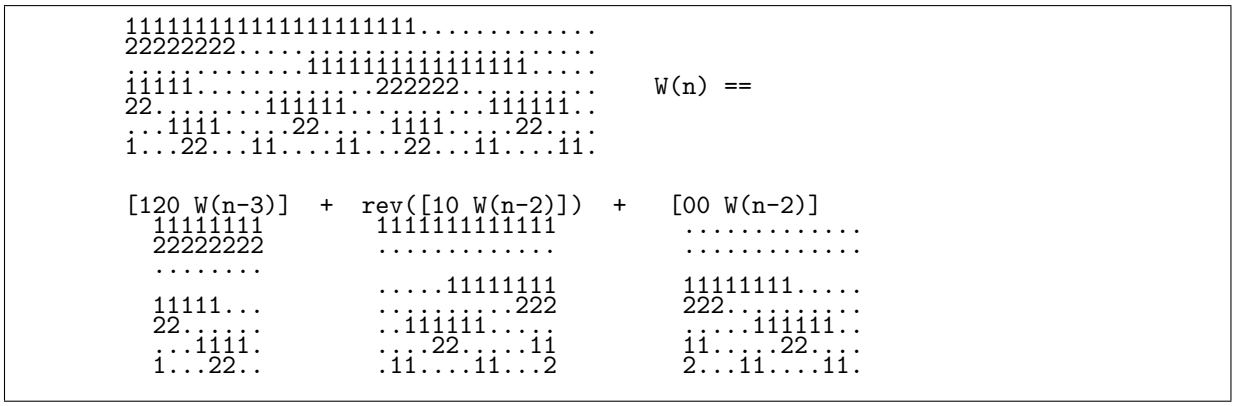

Figure 12.0-A: Obtaining a Gray code by a sublist recursion.

The algorithms are given as *list recursions*. For example, write $W(n)$ for the list of n-digit words (of a certain type), write $W^{\mathbf{R}}(n)$ for the reversed list, and $[x\,.\,W(n)]$ for the list with the word x prepended at each word. The recursion for a Gray code is

$$W(n) \;=\; \begin{matrix}[0\,0\,.\,W(n-2)\;\;] \\ [1\,0\,.\,W^{\mathbf{R}}(n-2)] \\ [1\,2\,0\,.\,W(n-3)]\end{matrix} \tag{12.0-1}$$

A relation like this always implies a backward version which is obtained by reversing the order of the sublists on the right hand side and additionally reversing each sublist

$$W^{\mathbf{R}}(n) \;=\; \begin{matrix}[1\,2\,0\,.\,W^{\mathbf{R}}(n-3)] \\ [1\,0\,.\,W(n-2)\;\;\;] \\ [0\,0\,.\,W^{\mathbf{R}}(n-2)\;\;] \end{matrix} \tag{12.0-2}$$

The construction is illustrated in figure 12.0-A. An implementation of the algorithm is [FXT: comb/fib-alt-gray-demo.cc]:

```
1   void X_rec(ulong d, bool z)
2   {
3       if ( d>=n )
4       {
5           if ( d<=n+1 )  // avoid duplicates
6           {
7               visit();
8           }
9       }
10      else
11      {
12          if ( z )
13          {
14              rv[d]=0;  rv[d+1]=0;  X_rec(d+2, z);
15              rv[d]=1;  rv[d+1]=0;  X_rec(d+2,  ! z);
16              rv[d]=1;  rv[d+1]=2;  rv[d+2]=0;  X_rec(d+3, z);
17          }
18          else
19          {
20              rv[d]=1;  rv[d+1]=2;  rv[d+2]=0;  X_rec(d+3, z);
21              rv[d]=1;  rv[d+1]=0;  X_rec(d+2,  ! z);
22              rv[d]=0;  rv[d+1]=0;  X_rec(d+2, z);
23          }
24      }
25  }
```

The initial call is `X_rec(0, 0);`. The parameter `z` determines whether the list is generated in forward or backward order. No optimizations are made as these tend to obscure the idea. Here we could omit one statement `rv[d]=1;`, replace the arguments `z` and `!z` in the recursive calls by constants, or create an iterative version.

The number $w(n)$ of words $W(n)$ is determined by a recursion (and some initial values $w(n)$) that can be obtained by counting the size of the list on both sides of the recursion relation 12.0-1:

$$w(n) \quad = \quad 2\,w(n-2) + w(n-3) \tag{12.0-3}$$

One can typically set $w(0) = 1$, there is one empty list and this satisfies all conditions. The numbers $w(n)$ are in fact the Fibonacci numbers.

12.1 Fibonacci words

A recursive routine to generate the Fibonacci words (binary words not containing two consecutive ones) can be given as follows:

```
1   ulong n;    // number of bits in words
2   ulong *rv;  // bits of the word
3
4   void fib_rec(ulong d)
5   {
6       if ( d>=n )  visit();
7       else
8       {
9            rv[d]=0;  fib_rec(d+1);
10           rv[d]=1;  rv[d+1]=0;  fib_rec(d+2);
11      }
12  }
```

We allocate one extra element to avoid if-statements in the code:

```
int main()
{
    n = 7;
    rv = new ulong[n+1];
    fib_rec(0);
    return 0;
}
```

The output (assuming `visit()` simply prints the array) is given in the left of figure 12.1-A. Note that with the n-bit Fibonacci Gray code the number of ones in the first and last, second and second-last, etc. tracks are equal. Thereby the sequence of reversed words is also a Fibonacci Gray code. A simple modification of the routine generates a Gray code through the Fibonacci words [FXT: comb/fibgray-rec-demo.cc]:

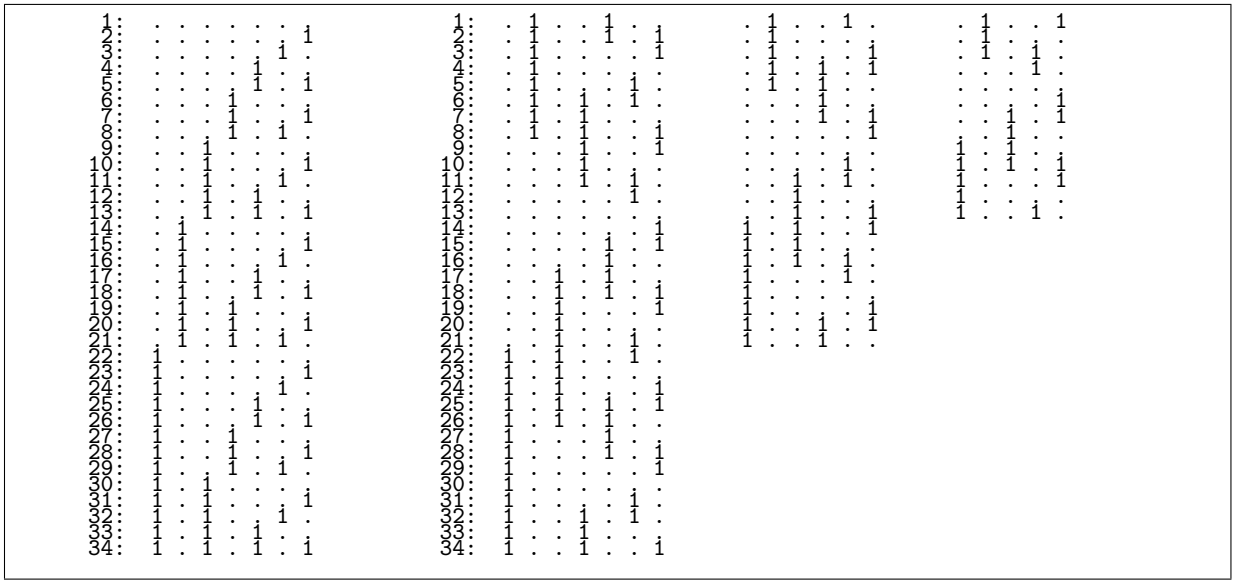

Figure 12.1-A: The first 34 Fibonacci words in counting order (left), and Gray codes through the first 34, 21, and 13 Fibonacci words (right). Dots are used for zeros.

```
1   void fib_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           z = !z;  // change direction for Gray code
7           if ( z )
8           {
9               rv[d]=0;  fib_rec(d+1, z);
10              rv[d]=1;  rv[d+1]=0;  fib_rec(d+2, z);
11          }
12          else
13          {
14              rv[d]=1;  rv[d+1]=0;  fib_rec(d+2, z);
15              rv[d]=0;  fib_rec(d+1, z);
16          }
17      }
18  }
```

The variable z controls the direction in the recursion, it is changed unconditionally with each step. The `if-else` blocks can be merged into

```
1           rv[d]=!z;  rv[d+1]= z;  fib_rec(d+1+!z, z);
2           rv[d]= z;  rv[d+1]=!z;  fib_rec(d+1+ z, z);
```

The algorithm is CAT (constant amortized time) and fast in practice, about 80 million objects are generated per second. A bit-level algorithm is given in section 1.30.2 on page 79.

The algorithm for the list of the length-n Fibonacci words $F(n)$ can be given as a recursion:

$$F(n) \;=\; \begin{matrix}[1\,0\,.\,F^{\mathbf{R}}(n-2)] \\ [0\,.\,F^{\mathbf{R}}(n-1)\;\;]\end{matrix} \tag{12.1-1}$$

Merging two steps we find

$$F(n) \;=\; \begin{matrix}[1\,0\,0\,.\,F(n-3)\;\;] \\ [1\,0\,1\,0\,.\,F(n-4)] \\ [0\,0\,.\,F(n-2)\;\;\;] \\ [0\,1\,0\,.\,F(n-3)\;\;]\end{matrix} \tag{12.1-2}$$

12.2 Generalized Fibonacci words

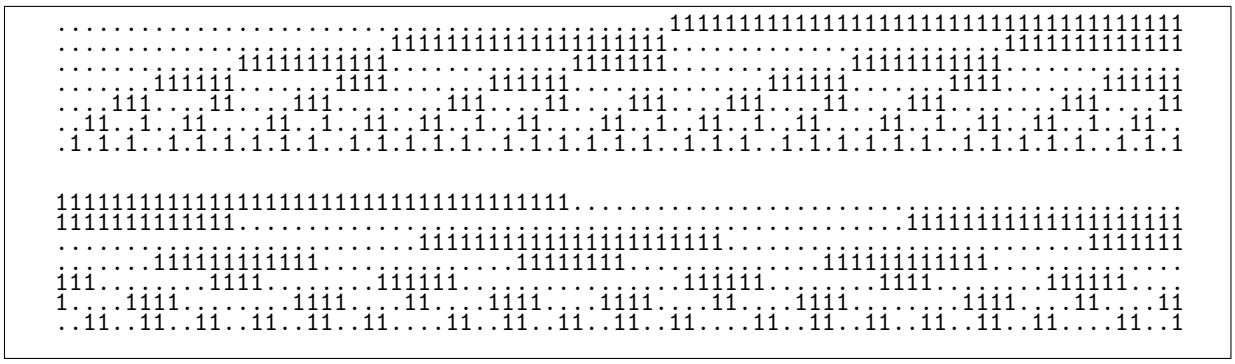

Figure 12.2-A: The 7-bit binary words with maximal 2 successive ones in lexicographic (top) and minimal-change (bottom) order. Dots denote zeros.

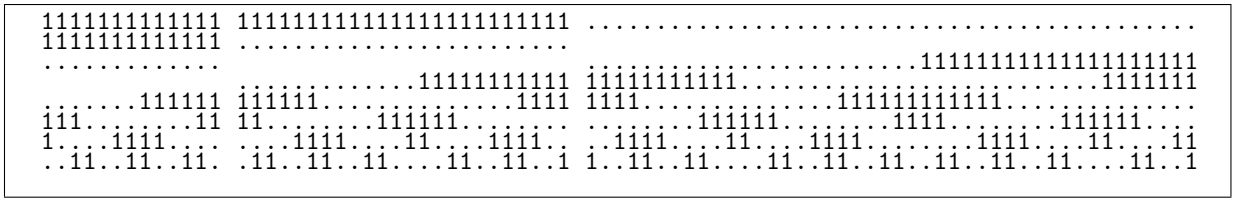

Figure 12.2-B: Recursive structure for the 7-bit binary words with maximal 2 successive ones.

We generalize the Fibonacci words by allowing a fixed maximum value r of successive ones in a binary word. The Fibonacci words correspond to $r = 1$. Figure 12.2-A shows the 7-bit words with $r = 2$. The method to generate a Gray code for these words is a generalization of the recursion for the Fibonacci words. Write $L_r(n)$ for the list of n-bit words with at most r successive ones, then the recursive structure for the Gray code is

$$L_r(n) = \begin{matrix} [0\,.\,L_r^{\mathbf{R}}(n-1) \quad\quad\quad\

```
15              {
16                  rv[d+k] = 0;
17                  maxrep_rec(d+1+k, z);
18                  rv[d+k] = 1;
19              }
20          }
21          else
22          {
23              // words: ... 1110, 110, 10, 0
24              for (long k=0; k<km; ++k)  rv[d+k] = 1;
25              for (long k=km; k>=0; --k)
26              {
27                  rv[d+k] = 0;
28                  maxrep_rec(d+1+k, z);
29              }
30          }
31      }
32  }
```

Figure 12.2-C shows the 5-bit Gray codes for $r \in \{1, 2, 3, 4, 5\}$. Observe that all sequences are subsequences of the leftmost column.

```
          r = 5       r = 4       r = 3       r = 2       r = 1
 1:   1 1 1 1 1   1 1 1 1 .   1 1 1 . .   1 1 . . 1   1 . . 1 .
 2:   1 1 1 1 .   1 1 1 . .   1 1 1 . 1   1 1 . . .   1 . . . .
 3:   1 1 1 . .   1 1 1 . 1   1 1 . . 1   1 1 . 1 .   1 . . . 1
 4:   1 1 1 . 1   1 1 . . 1   1 1 . . .   1 1 . 1 1   1 . 1 . 1
 5:   1 1 . . 1   1 1 . . .   1 1 . 1 .   1 . . 1 1   1 . 1 . .
 6:   1 1 . . .   1 1 . 1 .   1 1 . 1 1   1 . . 1 .   . . 1 . .
 7:   1 1 . 1 .   1 1 . 1 1   1 . . 1 1   1 . . . .   . . 1 . 1
 8:   1 1 . 1 1   1 . . 1 1   1 . . 1 .   1 . . . 1   . . . . 1
 9:   1 . . 1 1   1 . . 1 .   1 . . . .   1 . 1 . 1   . . . . .
10:   1 . . 1 .   1 . . . .   1 . . . 1   1 . 1 . .   . . . 1 .
11:   1 . . . .   1 . . . 1   1 . 1 . 1   1 . 1 1 .   . 1 . 1 .
12:   1 . . . 1   1 . 1 . 1   1 . 1 . .   . . 1 1 .   . 1 . . .
13:   1 . 1 . 1   1 . 1 . .   1 . 1 1 .   . . 1 . .   . 1 . . 1
14:   1 . 1 . .   1 . 1 1 .   1 . 1 1 1   . . 1 . 1
15:   1 . 1 1 .   1 . 1 1 1   . . 1 1 1   . . . . 1
16:   1 . 1 1 1   . . 1 1 1   . . 1 1 .   . . . . .
17:   . . 1 1 1   . . 1 1 .   . . 1 . .   . . . 1 .
18:   . . 1 1 .   . . 1 . .   . . 1 . 1   . . . 1 1
19:   . . 1 . .   . . 1 . 1   . . . . 1   . 1 . 1 1
20:   . . 1 . 1   . . . . 1   . . . . .   . 1 . 1 .
21:   . . . . 1   . . . . .   . . . 1 .   . 1 . . .
22:   . . . . .   . . . 1 .   . . . 1 1   . 1 . . 1
23:   . . . 1 .   . . . 1 1   . 1 . 1 1   . 1 1 . 1
24:   . . . 1 1   . 1 . 1 1   . 1 . 1 .   . 1 1 . .
25:   . 1 . 1 1   . 1 . 1 .   . 1 . . .
26:   . 1 . 1 .   . 1 . . .   . 1 . . 1
27:   . 1 . . .   . 1 . . 1   . 1 1 . 1
28:   . 1 . . 1   . 1 1 . 1   . 1 1 . .
29:   . 1 1 . 1   . 1 1 . .   . 1 1 1 .
30:   . 1 1 . .   . 1 1 1 .
31:   . 1 1 1 .   . 1 1 1 1
32:   . 1 1 1 1
```

Figure 12.2-C: Gray codes of the 5-bit binary words with maximal r successive ones. The leftmost column is the complement of the Gray code of all binary words, the rightmost column is the Gray code for the Fibonacci words.

Let $w_r(n)$ be the number of n-bit words $W_r(n)$ with $\leq r$ successive ones. Taking the length of the lists on both sides of relation 12.2-1 we obtain the recursion

$$w_r(n) \quad = \quad \sum_{j=0}^{r} w_r(n-1-j) \tag{12.2-2}$$

where we set $w_r(n) = 2^k$ for $0 \leq n \leq r$. The sequences for $r \leq 5$ start as

n:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
r=1:	1	2	3	5	8	13	21	34	55	89	144	233	377	610	987	1597
r=2:	1	2	4	7	13	24	44	81	149	274	504	927	1705	3136	5768	10609
r=3:	1	2	4	8	15	29	56	108	208	401	773	1490	2872	5536	10671	20569
r=4:	1	2	4	8	16	31	61	120	236	464	912	1793	3525	6930	13624	26784
r=5:	1	2	4	8	16	32	63	125	248	492	976	1936	3840	7617	15109	29970

For $r = 1$ we obtain the *Fibonacci numbers*, entry A000045 of [240]; For $r = 2$ the *tribonacci numbers*, entry A000073; for $r = 3$ the *tetranacci numbers*, entry A000078; for $r = 4$ the *pentanacci numbers*, entry A001591; for $r = 5$ the *hexanacci numbers*, entry A001592. The variant of the Fibonacci sequence where each number is the sum of its k predecessors is also called *Fibonacci k-step sequence*. The generating function for $w_r(n)$ is

$$\sum_{n=0}^{\infty} w_r(n)\, x^n \quad = \quad \frac{\sum_{k=0}^{r} x^k}{1 - \sum_{k=1}^{r+1} x^k} \tag{12.2-3}$$

Alternative Gray code for words without substrings 111 ($r = 2$)

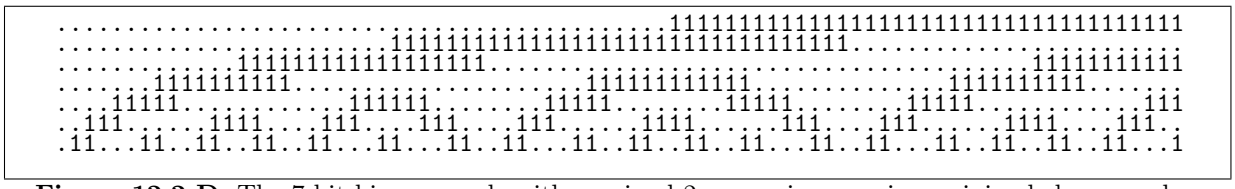

Figure 12.2-D: The 7-bit binary words with maximal 2 successive ones in a minimal-change order.

The list recursion for the Gray code for binary words without substrings 111 is the special case $r = 2$ of relation 12.2-1 on page 298:

$$L_2(n) \quad = \quad \begin{matrix} [1\,1\,0\,.\,L_2^{\mathbf{R}}(n-3)] \\ [1\,0\,.\,L_2^{\mathbf{R}}(n-2)\ \ ] \\ [0\,.\,L_2^{\mathbf{R}}(n-1)\ \ \ ] \end{matrix} \tag{12.2-4}$$

A different Gray code is obtained via

$$L'_2(n) \quad = \quad \begin{matrix} [1\,0\,.\,L'_2(n-2)\ \ \ ] \\ [1\,1\,0\,.\,L'^{\mathbf{R}}_2(n-3)] \\ [0\,.\,L'_2(n-1)\ \ \ ] \end{matrix} \tag{12.2-5}$$

The ordering is shown in figure 12.2-D, it was created with the program [FXT: comb/no111-gray-demo.cc].

Alternative Gray code for words without substrings 1111 ($r = 3$)

A list recursion for an alternative Gray code for binary words without substrings 1111 ($r = 3$) is

$$L'_3(n) \quad = \quad \begin{matrix} [1\,1\,0\,.\,L'^{\mathbf{R}}_3(n-3)\ \ ] \\ [0\,.\,L'^{\mathbf{R}}_3(n-1)\ \ \ \ ] \\ [1\,1\,1\,0\,.\,L'^{\mathbf{R}}_3(n-4)] \\ [1\,0\,.\,L'^{\mathbf{R}}_3(n-2)\ \ \ ] \end{matrix} \tag{12.2-6}$$

The ordering is shown in figure 12.2-E, it was created with the program [FXT: comb/no1111-gray-demo.cc]. For all odd $r \geq 3$ a Gray code can be obtained by a list recursion where the prefixes with an even number of ones are followed by those with an odd number of ones. For example, with $r = 5$ the recursion is

$$L'_5(n) \quad = \quad \begin{matrix} [1\,1\,1\,1\,0\,.\,L'^{\mathbf{R}}_5(n-7)\ \ ] \\ [1\,1\,0\,.\,L'^{\mathbf{R}}_5(n-3)\ \ \ \ ] \\ [0\,.\,L'^{\mathbf{R}}_5(n-1)\ \ \ \ ] \\ [1\,1\,1\,1\,1\,0\,.\,L'^{\mathbf{R}}_5(n-6)] \\ [1\,1\,1\,0\,.\,L'^{\mathbf{R}}_5(n-4)\ \ ] \\ [1\,0\,.\,L'^{\mathbf{R}}_5(n-2)\ \ \ \ ] \end{matrix} \tag{12.2-7}$$

```
  1:  1.11.1.      28:  1.1..1.      55:  ..111.1      82:  .1.1.1.
  2:  1.11...      29:  1.1.11.      56:  ...11.1      83:  .1.111.
  3:  1.11..1      30:  111.11.      57:  ...11..      84:  .1..11.
  4:  1.11.11      31:  111..1.      58:  ....1.,      85:  .1...1.
  5:  1..1.11      32:  111....      59:  ....1.1      86:  .1.....
  6:  1..1..1      33:  111...1      60:  ....111      87:  .1....1
  7:  1..1...      34:  111..11      61:  .....11      88:  .1...11
  8:  1..1.1.      35:  111.111      62:  ......1      89:  .1..111
  9:  1..111.      36:  111.1.1      63:  .......      90:  .1..1.1
 10:  1...11.      37:  111.1..      64:  .....1.      91:  .1..1..
 11:  1....1.      38:  .11.1.,      65:  ....11.      92:  .1.11.,
 12:  1......      39:  .11.1.1      66:  ...111.      93:  .1.11.1
 13:  1.....1      40:  .11.111      67:  ...1.1.      94:  11.11.1
 14:  1....11      41:  .11..11      68:  ...1...      95:  11.11..
 15:  1...111      42:  .11...1      69:  ...1..1      96:  11..1..
 16:  1...1.1      43:  .11....      70:  ...1.11      97:  11..1.1
 17:  1...1..      44:  .11..1.      71:  ..11.11      98:  11..111
 18:  1..11..      45:  .11.11.      72:  ..11..1      99:  11...11
 19:  1..11.1      46:  ..1.11.      73:  ..11...     100:  11....1
 20:  1.111.1      47:  ..1..1.      74:  ..11.1.     101:  11.....
 21:  1.111..      48:  ..1....      75:  .111.1.     102:  11...1.
 22:  1.1.1..      49:  ..1...1      76:  .111...     103:  11..11.
 23:  1.1.1.1      50:  ..1..11      77:  .111..1     104:  11.111.
 24:  1.1.111      51:  ..1.111      78:  .111.11     105:  11.1.1.
 25:  1.1..11      52:  ..1.1.1      79:  .1.1.11     106:  11.1..,
 26:  1.1...1      53:  ..1.1..      80:  .1.1..1     107:  11.1..1
 27:  1.1....      54:  ..111..      81:  .1.1...     108:  11.1.11
```

Figure 12.2-E: The 7-bit binary words with maximal 3 successive ones in a minimal-change order.

12.3 Digit x followed by at least x zeros

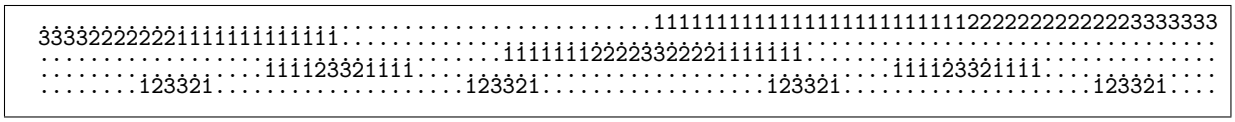

Figure 12.3-A: Gray code for the length-5 words with maximal digit 3 where a digit x is followed by at least x zeros. Dots denote zeros.

Figure 12.3-A shows a Gray code for the length-5 words with maximal digit 3 where a digit x is followed by at least x zeros. For the Gray code list $Z_r(n)$ of the length-n words with maximal digit r we have

$$Z_r(n) \;=\; \begin{bmatrix} [0\,.\,Z_r^{\mathbf{R}}(n-1) &] \\ [1\,0\,.\,Z_r^{\mathbf{R}}(n-2) &] \\ [2\,0\,0\,.\,Z_r^{\mathbf{R}}(n-3) &] \\ [3\,0\,0\,0\,.\,Z_r^{\mathbf{R}}(n-4) &] \\ [\quad\vdots &] \\ [r\,0^r\,.\,Z_r^{\mathbf{R}}(n-r-1) &] \end{bmatrix} \tag{12.3-1}$$

A implementation is [FXT: comb/gexz-gray-demo.cc]:

```
ulong n;    // number of digits in words
ulong *rv;  // digits of the word
ulong mr;   // radix== mr+1
```

```
void gexz_rec(ulong d, bool z)
{
    if ( d>=n )  visit();
    else
    {
        if ( z )
        {
            // words 0, 10, 200, 3000, 40000, ...
            ulong k = 0;
            do
            {
                rv[d]=k;
                for (ulong j=1; j<=k; ++j)   rv[d+j] = 0;
                gexz_rec(d+k+1, !z);
```

```
15              }
16              while ( ++k <= mr );
17          }
18          else
19          {
20              // words ..., 40000, 3000, 200, 10, 0
21              ulong k = mr + 1;
22              do
23              {
24                  --k;
25                  rv[d]=k;
26                  for (ulong j=1; j<=k; ++j)   rv[d+j] = 0;
27                  gexz_rec(d+k+1, !z);
28              }
29              while ( k != 0 );
30          }
31      }
32  }
```

Let $z_r(n)$ be the number of n-bit words $Z_r(n)$, then

$$z_r(n) \quad = \quad \sum_{j=1}^{r+1} z_r(n-j) \tag{12.3-2}$$

where we set $z_r(n) = 1$ for $n \leq 0$. The sequences for $r \leq 5$ start as

n:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
r=1:	1	2	3	5	8	13	21	34	55	89	144	233	377	610	987	1597
r=2:	1	3	5	9	17	31	57	105	193	355	653	1201	2209	4063	7473	13745
r=3:	1	4	7	13	25	49	94	181	349	673	1297	2500	4819	9289	17905	34513
r=4:	1	5	9	17	33	65	129	253	497	977	1921	3777	7425	14597	28697	56417
r=5:	1	6	11	21	41	81	161	321	636	1261	2501	4961	9841	19521	38721	76806

For $r = 1$ we obtain the *Fibonacci numbers*, entry A000045 of [240]; For $r = 2$ the *tribonacci numbers*, entry A000213; for $r = 3$ the *tetranacci numbers*, entry A000288; for $r = 4$ the *pentanacci numbers*, entry A000322; for $r = 5$ the *hexanacci numbers*, entry A000383. Note that the sequences for $r \geq 2$ are different from those obtained via relation 12.2-2 on page 299.

12.4 Generalized Pell words

12.4.1 Gray code for Pell words

```
  ..................................................1111111111111111111111111111111111  2222222
  ..................11111111111111112222222.....................11111111111111111  .......
  .......1111111222.......1111111222.................1111111222.......111111122  1111222
  ...1112...1112......1112...1112......1112...1112...1112......1112...1112..  1112...
  .12.12..12.12..12.12.12..12.12..12.12.12..12.12..12.12..12.12.12..12.12..1  .12..12

  ..................................................1111111111111111111111111111111111  2222222
  ..................11111111111111112222222222222211111111111111111.........  .......
  .......1111111222222111111.....................................1111111222222111111  1111222
  ...11122111.............11122111.....11122111.....11122111.............111  1.....
  .1221....1221..1221..1221....1221..1221....1221..1221....1221..1221..1221.  221..12
```

Figure 12.4-A: Start and end of the lists of five digit Pell words in counting order (top) and Gray code order (bottom). The lowest row is the least significant digit, dots denote zeros.

A Gray code of the Pell words (ternary words without the substrings `"21"` and `"22"`) can be computed as follows:

```
1   ulong n;    // number of digits in words
2   ulong *rv;  // digits of the word
3   bool zq;    // order: 0==>Lex,  1==>Gray
```

```
1   void pell_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           if ( 0==z )
7           {
8               rv[d]=0;  pell_rec(d+1, z);
9               rv[d]=1;  pell_rec(d+1, zq^z);
10              rv[d]=2;  rv[d+1]=0;  pell_rec(d+2, z);
11          }
12          else
13          {
14              rv[d]=2;  rv[d+1]=0;  pell_rec(d+2, z);
15              rv[d]=1;  pell_rec(d+1, zq^z);
16              rv[d]=0;  pell_rec(d+1, z);
17          }
18      }
19  }
```

The global boolean variable `zq` controls whether the counting order or the Gray code is generated. The code is given in [FXT: comb/pellgray-rec-demo.cc]. Both orderings are shown in figure 12.4-A. About 110 million words per second are generated. The computation of a function whose power series coefficients are related to the Pell Gray code is described in section 36.12.3 on page 746.

12.4.2 Gray code for generalized Pell words

```
                                        .....11111111111111111111111111111
3333222222222222211111111111111.....................1111111111122222
........11112222332222111111.......1111222233222211111.......11112222332222
.123321..123321....123321..123321..123321....123321..123321..123321....123


  11111111111122222222222222222222222222222222222222223333333333333
  22222222333333332222222222222111111111111.........................
  1111.............11112222332222111111.......1111222233222211111....
  321..123321..123321..123321....123321..123321..123321....123321..123
```

Figure 12.4-B: Gray code for 4-digit radix-4 strings with no substring $3x$ with $x \neq 0$.

A generalization of the Pell words are the radix-$(r+1)$ strings where the substring rx with $x \neq 0$ is forbidden. Let $P_r(n)$ the length-n words, a gray code for these strings can be generated by the recursion

$$P_r(n) = \begin{bmatrix} [0\,.\,P_r(n-1)] \\ [1\,.\,P_r^{\mathbf{R}}(n-1)] \\ [2\,.\,P_r(n-1)] \\ [3\,.\,P_r^{\mathbf{R}}(n-1)] \\ \vdots \\ [(r-1)\,.\,P_r^{\mathbf{R}}(n-1)] \\ [(r)\,0\,.\,P_r(n-2)] \end{bmatrix} \tag{12.4-1a}$$

if r is even, and

$$P_r(n) = \begin{bmatrix} [0\,.\,P_r^{\mathbf{R}}(n-1)] \\ [1\,.\,P_r(n-1)] \\ [2\,.\,P_r^{\mathbf{R}}(n-1)] \\ [3\,.\,P_r(n-1)] \\ \vdots \\ [(r-1)\,.\,P_r(n-1)] \\ [(r)\,0\,.\,P_r^{\mathbf{R}}(n-2)] \end{bmatrix} \tag{12.4-1b}$$

if r is odd. Figure 12.4-B shows a Gray code for the 4-digit strings with $r = 3$. An implementation of the algorithm is [FXT: comb/pellgen-gray-demo.cc]:

```
1   ulong n;    // number of digits in words
2   ulong *rv;  // digits of the word (radix r+1)
3   long r;     // Forbidden substrings are [r, x] where x!=0
```

```
1   void pellgen_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           const bool p = r & 1;  // parity of r
7           rv[d] = 0;
8           if ( z )
9           {
10              for (long k=0; k<r; ++k)  { rv[d] = k;  pellgen_rec(d+1, z ^ p ^ (k&1)); }
11              { rv[d] = r;  rv[d+1] = 0;  pellgen_rec(d+2, p ^ z); }
12          }
13          else
14          {
15              { rv[d] = r;  rv[d+1] = 0;  pellgen_rec(d+2, p ^ z); }
16              for (long k=r-1; k>=0; --k)  { rv[d] = k;  pellgen_rec(d+1, z ^ p ^ (k&1)); }
17          }
18      }
19  }
```

With $r = 1$ we again obtain the Gray code for Fibonacci words. Taking the number $p_r(n)$ of words $P_r(n)$ on both sides of relations 12.4-1a and 12.4-1b we find

$$p_r(n) \quad = \quad r\,p_r(n) + p_r(n-2) \tag{12.4-2}$$

where $p_r(0) = 1$ and $p_r(1) = r + 1$. For $r \le 5$ the sequences start as

n:	0	1	2	3	4	5	6	7	8	9	10	11
r=1:	1	2	3	5	8	13	21	34	55	89	144	233
r=2:	1	3	7	17	41	99	239	577	1393	3363	8119	19601
r=3:	1	4	13	43	142	469	1549	5116	16897	55807	184318	608761
r=4:	1	5	21	89	377	1597	6765	28657	121393	514229	2178309	9227465
r=5:	1	6	31	161	836	4341	22541	117046	607771	3155901	16387276	85092281

The sequences are the following entries of [240]: $r = 1$: A000045; $r = 2$: A001333; $r = 3$: A003688; $r = 4$: A015448; $r = 5$: A015449. The generating function for $p_r(n)$ is

$$\sum_{n=0}^{\infty} p_r(n)\,x^n \quad = \quad \frac{1+x}{1-r\,x-x^2} \tag{12.4-3}$$

12.5 Sparse signed binary words

```
.....................................................MMMMMMMMMMMMMMMMMMMMPPPPPPPPPPPPPPPPPPPPP
PPPPPPPPPPPPMMMMMMMMMMM.....................................................................
.....................................MMMMMPPPPPPPPPPPMMMMM.....................MMMMMPPPPP
PPPMMM.........MMMPPPPPPMMM.............................MMMPPPPPPMMM.............MPPM.....MP
PM..MPPM......MPPM..MPPM..MPPM......MPPM......MPPM......MPPM..MPPM..MPPM......MPPM...
```

Figure 12.5-A: A Gray code through the 85 sparse 6-bit signed binary words. Dots are used for zeros, the symbols '`P`' and '`M`' denote $+1$ and -1, respectively.

Figure 12.5-A shows a minimal-change order (Gray code) for the sparse signed binary words (nonadjacent form (NAF), see section 1.25 on page 64). Note that we allow a digit to switch between $+1$ and -1. If all words with any positive digit ('`P`') are omitted then we obtain the Gray code for Fibonacci words given in section 12.1 on page 296.

A recursive routine for the generation of the Gray code is given in [FXT: comb/naf-gray-rec-demo.cc]:

```
1   ulong n;   // number of digits of the string
2   int *rv;   // the string
```

```
1   void sb_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           if ( 0==z )
7           {
8               rv[d]=0;  sb_rec(d+1, 1);
9               rv[d]=-1;  rv[d+1]=0;  sb_rec(d+2, 1);
10              rv[d]=+1;  rv[d+1]=0;  sb_rec(d+2, 0);
11          }
12          else
13          {
14              rv[d]=+1;  rv[d+1]=0;  sb_rec(d+2, 1);
15              rv[d]=-1;  rv[d+1]=0;  sb_rec(d+2, 0);
16              rv[d]=0;  sb_rec(d+1, 0);
17          }
18      }
19  }
```

About 120 million words per second are generated.

Let $S(n)$ be the number of n-digit sparse signed binary numbers (of both signs), and $P(n)$ be the number of positive n-digit sparse signed binary numbers, then

```
   n: 0 1 2  3  4  5  6  7   8   9   10   11   12   13    14    15    16
S(n): 1 3 5 11 21 43 85 171 341 683 1365 2731 5461 10923 21845 43691 87381
P(n): 1 2 3  6 11 22 43 86  171 342  683 1366 2731  5462 10923 21846 43691
```

The sequence $S(n)$ is entry A001045 of [240], the sequence $P(n)$ is entry A005578. We have (with $e := n \bmod 2$)

$$S(n) = \frac{2^{n+2}-1+2\,e}{3} = 2\,S(n-1)-1+2\,e \tag{12.5-1a}$$

$$= S(n-1)+2\,S(n-2) = 3\,S(n-2)+2\,S(n-3) = 2\,P(n)-1 \tag{12.5-1b}$$

$$P(n) = \frac{2^{n+1}+1+e}{3} = 2\,P(n-1)-1-e = S(n-1)+e \tag{12.5-1c}$$

$$= P(n-1)+S(n-2) = P(n-2)+S(n-2)+S(n-3) \tag{12.5-1d}$$

$$= S(n-2)+S(n-3)+S(n-4)+\ldots+S(2)+S(1)+3 \tag{12.5-1e}$$

$$= 2\,P(n-1)+P(n-2)-2\,P(n-3) \tag{12.5-1f}$$

Almost Gray code for positive words *

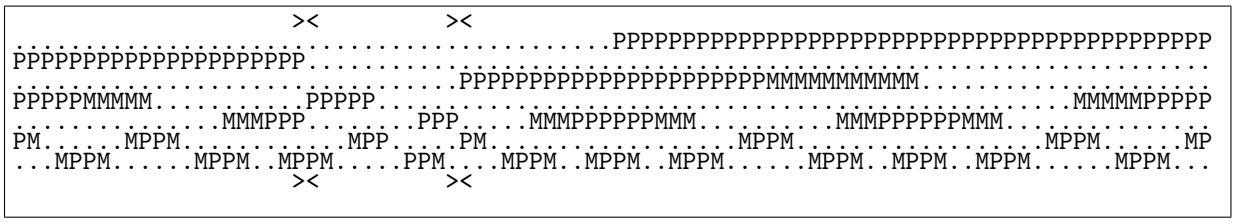

Figure 12.5-B: An ordering of the 86 sparse 7-bit positive signed binary words that is almost a Gray code. The transitions that are not minimal are marked with '><'. Dots denote zeros.

If we start with the following routine that calls `sb_rec()` only after a one has been inserted, we obtain an ordering of the positive numbers:

```
1   void pos_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           if ( 0==z )
7           {
8               rv[d]=0;  pos_rec(d+1, 1);
9               rv[d]=+1;  rv[d+1]=0;  sb_rec(d+2, 1);
```

```
10          }
11          else
12          {
13              rv[d]=+1;  rv[d+1]=0;  sb_rec(d+2, 0);
14              rv[d]=0;  pos_rec(d+1, 0);
15          }
16      }
17  }
```

The ordering obtained with n-digit words is a Gray code, except for $n-4$ transitions. An ordering with only about $n/2$ non-Gray transitions is obtained by the more complicated recursion [FXT: comb/naf-pos-rec-demo.cc]:

```
1   void pos_AAA(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           if ( 0==z )
7           {
8               rv[d]=+1;  rv[d+1]=0;  sb_rec(d+2, 0);  // 0
9               rv[d]=0;  pos_AAA(d+1, 1);  // 1
10          }
11          else
12          {
13              rv[d]=0;  pos_BBB(d+1, 0);  // 0
14              rv[d]=+1;  rv[d+1]=0;  sb_rec(d+2, 1);  // 1
15          }
16      }
17  }
```

```
1   void pos_BBB(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           if ( 0==z )
7           {
8               rv[d]=+1;  rv[d+1]=0;  sb_rec(d+2, 1);  // 1
9               rv[d]=0;  pos_BBB(d+1, 1);  // 1
10          }
11          else
12          {
13              rv[d]=0;  pos_AAA(d+1, 0);  // 0
14              rv[d]=+1;  rv[d+1]=0;  sb_rec(d+2, 0);  // 0
15          }
16      }
17  }
```

The initial call is `pos_AAA(0,0)`. The result for $n = 7$ is shown in figure 12.5-B. We list the number N of non-Gray transitions and the number of digit changes X in excess of a Gray code for $n \le 30$:

```
n: 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30
N: 0 0 0 0 1 2 2 2 3  4  4  4  5  6  6  6  7  8  8  8  9 10 10 10 11 12 12 12 13 14
X: 0 0 0 0 1 3 4 4 5  7  8  8  9 11 12 12 13 15 16 16 17 19 20 20 21 23 24 24 25 27
```

12.6 Strings with no two successive nonzero digits

A Gray code for the length-n strings with radix $(r+1)$ and no two successive nonzero digits is obtained by the recursion for the list $D_r(n)$:

$$D_r(n) \;=\; \begin{matrix} [\;0\,.\,D_r^{\mathbf{R}}(n-1)] \\ [1\,0\,.\,D_r^{\mathbf{R}}(n-1)] \\ [2\,0\,.\,D_r(n-1)\,] \\ [3\,0\,.\,D_r^{\mathbf{R}}(n-1)] \\ [4\,0\,.\,D_r(n-1)\,] \\ [5\,0\,.\,D_r^{\mathbf{R}}(n-1)] \\ [\quad\vdots\qquad\quad] \end{matrix} \tag{12.6-1}$$

An implementation is [FXT: comb/ntnz-gray-demo.cc]:

```
 1:  .3..3      26:  ....1      51:  1.1.2      76:  2.3.2
 2:  .3..2      27:  ....2      52:  1.1.3      77:  2.3.1
 3:  .3..1      28:  ....3      53:  1...3      78:  2.3..
 4:  .3...      29:  ..1.3      54:  1...2      79:  3.3..
 5:  .3.1.      30:  ..1.2      55:  1...1      80:  3.3.1
 6:  .3.2.      31:  ..1.1      56:  1....      81:  3.3.2
 7:  .3.3.      32:  ..1..      57:  1..1.      82:  3.3.3
 8:  .2.3.      33:  ..2..      58:  1..2.      83:  3.2.3
 9:  .2.2.      34:  ..2.1      59:  1..3.      84:  3.2.2
10:  .2.1.      35:  ..2.2      60:  2..3.      85:  3.2.1
11:  .2...      36:  ..2.3      61:  2..2.      86:  3.2..
12:  .2..1      37:  ..3.3      62:  2..1.      87:  3.1..
13:  .2..2      38:  ..3.2      63:  2....      88:  3.1.1
14:  .2..3      39:  ..3.1      64:  2...1      89:  3.1.2
15:  .1..3      40:  ..3..      65:  2...2      90:  3.1.3
16:  .1..2      41:  1.3..      66:  2...3      91:  3...3
17:  .1..1      42:  1.3.1      67:  2.1.3      92:  3...2
18:  .1...      43:  1.3.2      68:  2.1.2      93:  3...1
19:  .1.1.      44:  1.3.3      69:  2.1.1      94:  3....
20:  .1.2.      45:  1.2.3      70:  2.1..      95:  3..1.
21:  .1.3.      46:  1.2.2      71:  2.2..      96:  3..2.
22:  ...3.      47:  1.2.1      72:  2.2.1      97:  3..3.
23:  ...2.      48:  1.2..      73:  2.2.2
24:  ...1.      49:  1.1..      74:  2.2.3
25:  .....      50:  1.1.1      75:  2.3.3
```

Figure 12.6-A: Gray code for the length-4 radix-4 strings with no two successive nonzero digits.

```
1   ulong n;     // length of strings
2   ulong *rv;   // digits of strings
3   ulong mr;    // max digit
```

```
1   void ntnz_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           if ( 0==z )
7           {
8               rv[d]=0;  ntnz_rec(d+1, 1);
9               for (ulong t=1; t<=mr; ++t)  { rv[d]=t;  rv[d+1]=0;  ntnz_rec(d+2, t&1); }
10          }
11          else
12          {
13              for (ulong t=mr; t>0; --t)  { rv[d]=t;  rv[d+1]=0;  ntnz_rec(d+2, !(t&1)); }
14              rv[d]=0;  ntnz_rec(d+1, 0);
15          }
16      }
17  }
```

Figure 12.6-A shows the Gray code for length-4, radix-4 ($r = 3$) strings. With $r = 2$ and upon replacing 1 with -1 and 2 with $+1$ we obtain the Gray code for the sparse binary words (figure 12.5-A on page 304). With $r = 1$ we again obtain the Gray code for the Fibonacci words.

Counting the elements on both sides of relation 12.6-1 we find that for the number $d_r(n)$ of strings in the list $D_r(n)$ we have

$$d_r(n) \quad = \quad d_r(n-1) + r\, d_r(n-2) \tag{12.6-2}$$

where $d_r(0) = 1$ and $d_r(1) = r + 1$. The sequences of these numbers start as

```
   n: 0 1  2  3  4   5   6    7    8     9    10     11     12      13      14
r=1: 1 2  3  5  8  13  21   34   55    89   144    233    377     610     987
r=2: 1 3  5 11 21  43  85  171  341   683  1365   2731   5461   10923   21845
r=3: 1 4  7 19 40  97 217  508 1159  2683  6160  14209  32689   75316  173383
r=4: 1 5  9 29 65 181 441 1165 2929  7589 19305  49661 126881  325525  833049
r=5: 1 6 11 41 96 301 781 2286 6191 17621 48576 136681 379561 1062966 2960771
```

These are the following entries of [240]: $r = 1$: A000045; $r = 2$: A001045; $r = 3$: A006130; $r = 4$: A006131; $r = 5$: A015440; $r = 6$: A015441; $r = 7$: A015442; $r = 8$: A015443. The generating function for $d_r(n)$ is

$$\sum_{n=0}^{\infty} d_r(n)\, x^n \quad = \quad \frac{1 + r\, x}{1 - x - r\, x^2} \tag{12.6-3}$$

12.7 Strings with no two successive zeros

```
 1: . 1 3    21: 1 1 .    41: 2 3 1        1: . 1 2 2    21: 1 2 1 1    41: 2 . 1 .
 2: . 1 2    22: 1 1 1    42: 2 3 .        2: . 1 2 1    22: 1 2 1 2    42: 2 . 2 .
 3: . 1 1    23: 1 1 2    43: 3 3 .        3: . 1 2 .    23: 1 2 . 2    43: 2 . 2 1
 4: . 1 .    24: 1 1 3    44: 3 3 1        4: . 1 1 .    24: 1 2 . 1    44: 2 . 2 2
 5: . 2 .    25: 1 . 3    45: 3 3 2        5: . 1 1 1    25: 1 1 . 1    45: 2 1 2 2
 6: . 2 1    26: 1 . 2    46: 3 3 3        6: . 1 1 2    26: 1 1 . 2    46: 2 1 2 1
 7: . 2 2    27: 1 . 1    47: 3 2 3        7: . 1 . 2    27: 1 1 1 2    47: 2 1 2 .
 8: . 2 3    28: 2 . 1    48: 3 2 2        8: . 1 . 1    28: 1 1 1 1    48: 2 1 1 .
 9: . 3 3    29: 2 . 2    49: 3 2 1        9: . 2 . 1    29: 1 1 1 .    49: 2 1 1 1
10: . 3 2    30: 2 . 3    50: 3 2 .       10: . 2 . 2    30: 1 1 2 .    50: 2 1 1 2
11: . 3 1    31: 2 1 3    51: 3 1 .       11: . 2 1 2    31: 1 1 2 1    51: 2 1 . 2
12: . 3 .    32: 2 1 2    52: 3 1 1       12: . 2 1 1    32: 1 1 2 2    52: 2 1 . 1
13: 1 3 .    33: 2 1 1    53: 3 1 2       13: . 2 1 .    33: 1 . 2 2    53: 2 2 . 1
14: 1 3 1    34: 2 1 .    54: 3 1 3       14: . 2 2 .    34: 1 . 2 1    54: 2 2 . 2
15: 1 3 2    35: 2 2 .    55: 3 . 3       15: . 2 2 1    35: 1 . 2 .    55: 2 2 1 2
16: 1 3 3    36: 2 2 1    56: 3 . 2       16: . 2 2 2    36: 1 . 1 .    56: 2 2 1 1
17: 1 2 3    37: 2 2 2    57: 3 . 1       17: 1 2 2 2    37: 1 . 1 1    57: 2 2 1 .
18: 1 2 2    38: 2 2 3                    18: 1 2 2 1    38: 1 . 1 2    58: 2 2 2 .
19: 1 2 1    39: 2 3 3                    19: 1 2 2 .    39: 2 . 1 2    59: 2 2 2 1
20: 1 2 .    40: 2 3 2                    20: 1 2 1 .    40: 2 . 1 1    60: 2 2 2 2
```

Figure 12.7-A: Gray codes for strings with no two successive zeros: length-3 radix-4 (left), and length-4 radix-3 (right). Dots denote zeros.

A Gray code for the length-n strings with radix $(r+1)$ and no two successive zeros (see figure 12.7-A) is obtained by the recursion for the list $Z_r(n)$ as follows:

$$Z_r(n) = \begin{matrix} [0\,1\,.\,Z_r(n-2)\,] \\ [0\,2\,.\,Z_r^{\mathbf{R}}(n-2)] \\ [0\,3\,.\,Z_r(n-2)\,] \\ [0\,4\,.\,Z_r^{\mathbf{R}}(n-2)] \\ [0\,5\,.\,Z_r(n-2)\,] \\ [\quad\vdots\qquad\quad] \\ [0\,r\,.\,Z_r^{\mathbf{R}}(n-2)] \\ [1\,.\,Z_r^{\mathbf{R}}(n-1)\;] \\ [2\,.\,Z_r(n-1)\;\;] \\ [3\,.\,Z_r^{\mathbf{R}}(n-1)\;] \\ [4\,.\,Z_r(n-1)\;\;] \\ [\quad\vdots\qquad\quad] \\ [r\,.\,Z_r^{\mathbf{R}}(n-1)\;] \end{matrix} \text{ for } r \text{ even}, \qquad Z_r(n) = \begin{matrix} [0\,1\,.\,Z_r^{\mathbf{R}}(n-2)] \\ [0\,2\,.\,Z_r(n-2)\,] \\ [0\,3\,.\,Z_r^{\mathbf{R}}(n-2)] \\ [0\,4\,.\,Z_r(n-2)\,] \\ [0\,5\,.\,Z_r^{\mathbf{R}}(n-2)] \\ [\quad\vdots\qquad\quad] \\ [0\,r\,.\,Z_r^{\mathbf{R}}(n-2)] \\ [1\,.\,Z_r(n-1)\;\;] \\ [2\,.\,Z_r^{\mathbf{R}}(n-1)\;] \\ [3\,.\,Z_r(n-1)\;\;] \\ [4\,.\,Z_r^{\mathbf{R}}(n-1)\;] \\ [\quad\vdots\qquad\quad] \\ [r\,.\,Z_r(n-1)\;\;] \end{matrix} \text{ for } r \text{ odd}. \qquad (12.7\text{-}1)$$

An implementation is given in [FXT: comb/ntz-gray-demo.cc]:

```
1   ulong n;    // number of digits in words
2   ulong *rv;  // digits of the word (radix r+1)
3   long r;     // Forbidden substrings are [r, x] where x!=0

1   void ntz_rec(ulong d, bool z)
2   {
3       if ( d>=n )  visit();
4       else
5       {
6           bool w = 0;  // r-parity:  w depends on z ...
7           if ( r&1 )  w = !z;   // ... if r odd
8
9           if ( z )
10          {
11              // words 0X:
12              rv[d] = 0;
13              if ( d+2<=n )
14              {
15                  for (long k=1; k<=r; ++k, w=!w)  { rv[d+1]=k;  ntz_rec(d+2, w); }
16              }
17              else
18              {
```

```
19                  ntz_rec(d+1, w);
20                  w = !w;
21              }
22
23              w ^= (r&1);  //  r-parity:  change direction if r odd
24
25              // words X:
26              for (long k=1; k<=r; ++k, w=!w)  { rv[d]=k;  ntz_rec(d+1, w); }
27          }
28          else
29          {
30              // words X:
31              for (long k=r; k>=1; --k, w=!w)  { rv[d]=k;  ntz_rec(d+1, w); }
32
33              w ^= (r&1);  //  r-parity:  change direction if r odd
34
35              // words 0X:
36              rv[d] = 0;
37              if ( d+2<=n )
38              {
39                  for (long k=r; k>=1; --k, w=!w)  { rv[d+1]=k;  ntz_rec(d+2, w); }
40              }
41              else
42              {
43                  ntz_rec(d+1, w);
44                  w = !w;
45              }
46          }
47      }
48  }
```

With $r = 1$ we obtain the complement of the minimal-change list of Fibonacci words. Let $z_r(n)$ be the number of words $W_r(n)$, we find

$$z_r(n) \quad = \quad r\, z_r(n-1) + r\, z_r(n-1) \tag{12.7-2}$$

where $z_r(0) = 1$ and $z_r(1) = r + 1$. The sequences for $r \le 5$ start

n:	0	1	2	3	4	5	6	7	8	9	10	11
r=1:	1	2	3	5	8	13	21	34	55	89	144	233
r=2:	1	3	8	22	60	164	448	1224	3344	9136	24960	68192
r=3:	1	4	15	57	216	819	3105	11772	44631	169209	641520	2432187
r=4:	1	5	24	116	560	2704	13056	63040	304384	1469696	7096320	34264064
r=5:	1	6	35	205	1200	7025	41125	240750	1409375	8250625	48300000	282753125

These (for $r \le 4$) are the following entries of [240]: $r = 1$: A000045; $r = 2$: A028859; $r = 3$: A125145; $r = 4$: A086347. The generating function for $z_r(n)$ is

$$\sum_{n=0}^{\infty} z_r(n)\, x^n \quad = \quad \frac{1+x}{1 - r\,x - r\,x^2} \tag{12.7-3}$$

12.8 Binary strings without substrings 1x1 or 1xy1

12.8.1 No substrings 1x1

A Gray code for binary strings with no substring 1x1 is shown in figure 12.8-A. The recursive structure for the list $V(n)$ of the n-bit words is

$$V(n) \quad = \quad \begin{matrix} [1\,0\,0\,.\,V(n-3) \quad\quad\;\;] \\ [1\,1\,0\,0\,.\,V^{\mathbf{R}}(n-4)] \\ [0\,.\,V(n-1) \quad\quad\;\;] \end{matrix} \tag{12.8-1}$$

The implied algorithm can be implemented as [FXT: comb/no1x1-gray-demo.cc]:

```
1   ulong n;    // number of bits in words
2   ulong *rv;  // bits of the word
```

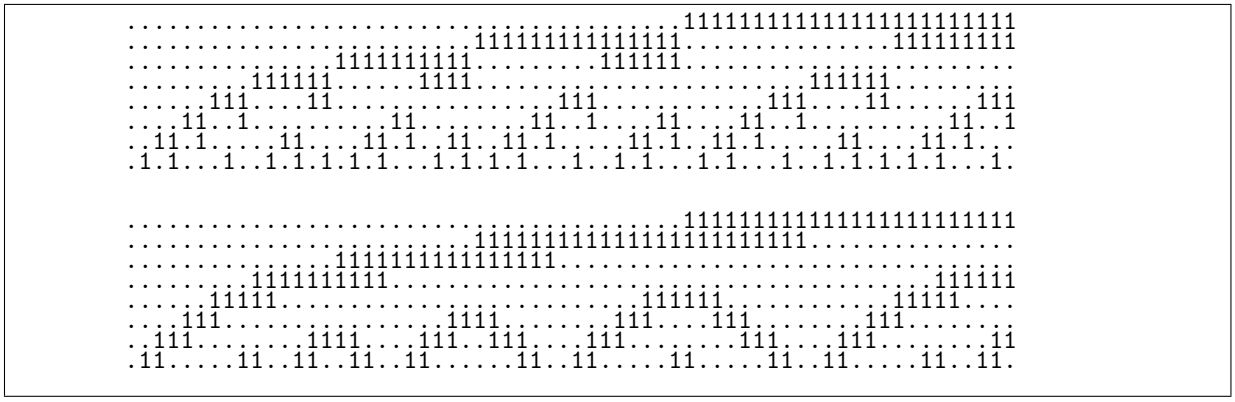

Figure 12.8-A: The length-8 binary strings with no substring 1x1 (where x is either 0 or 1): lex order (top) and minimal-change order (bottom). Dots denote zeros.

```
1   void no1x1_rec(ulong d, bool z)
2   {
3       if ( d>=n )  { if ( d<=n+2 )  visit(); }
4       else
5       {
6           if ( z )
7           {
8               rv[d]=1;  rv[d+1]=0;  rv[d+2]=0;  no1x1_rec(d+3, z);
9               rv[d]=1;  rv[d+1]=1;  rv[d+2]=0;  rv[d+3]=0;  no1x1_rec(d+4, !z);
10              rv[d]=0;  no1x1_rec(d+1, z);
11          }
12          else
13          {
14              rv[d]=0;  no1x1_rec(d+1, z);
15              rv[d]=1;  rv[d+1]=1;  rv[d+2]=0;  rv[d+3]=0;  no1x1_rec(d+4, !z);
16              rv[d]=1;  rv[d+1]=0;  rv[d+2]=0;  no1x1_rec(d+3, z);
17          }
18      }
19  }
```

The sequence of the numbers $v(n)$ of length-n strings starts as

```
n:     0 1 2 3 4  5  6  7  8   9  10  11  12  13   14   15   16   17
v(n):  1 2 4 6 9 15 25 40 64 104 169 273 441 714 1156 1870 3025 4895
```

This is entry A006498 of [240]. The recurrence relation is

$$v(n) \quad = \quad v(n-1) + v(n-3) + v(n-4) \tag{12.8-2}$$

The generating function is

$$\sum_{n=0}^{\infty} v(n)\, x^n \quad = \quad \frac{1 + x + 2\,x^2 + x^3}{1 - x - x^3 - x^4} \tag{12.8-3}$$

12.8.2 No substrings 1xy1

Figure 12.8-B shows a Gray code for binary words with no substring 1xy1. The recursion for the list of n-bit words $Y(n)$ is

$$Y(n) \quad = \quad \begin{array}{l} [1\,0\,0\,0\,.\,Y(n-4) \quad\quad\;\;] \\ [1\,0\,1\,0\,0\,0\,.\,Y^{\mathbf{R}}(n-6)] \\ [1\,1\,1\,0\,0\,0\,.\,Y(n-6) \;\;] \\ [1\,1\,0\,0\,0\,.\,Y^{\mathbf{R}}(n-5) \;\;] \\ [0\,.\,Y(n-1) \quad\quad\;\;] \end{array} \tag{12.8-4}$$

An implementation is given in [FXT: comb/no1xy1-gray-demo.cc]:

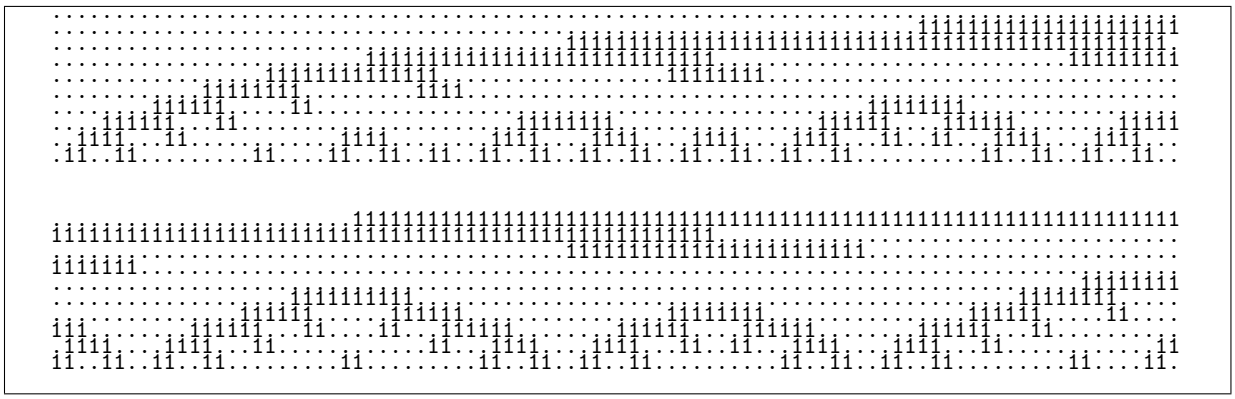

Figure 12.8-B: The length-10 binary strings with no substring 1xy1 (where x and y are either 0 or 1) in minimal-change order. Dots denote zeros.

```
1   void Y_rec(long p1, long p2, bool z)
2   {
3       if ( p1>p2 )  { visit();  return; }
4
5   #define S1(a)  rv[p1+0]=a
6   #define S2(a,b)  S1(a); rv[p1+1]=b;
7   #define S3(a,b,c)  S2(a,b); rv[p1+2]=c;
8   #define S4(a,b,c,d)  S3(a,b,c); rv[p1+3]=d;
9   #define S5(a,b,c,d,e)  S4(a,b,c,d); rv[p1+4]=e;
10  #define S6(a,b,c,d,e,f)  S5(a,b,c,d,e); rv[p1+5]=f;
11
12      long d = p2 - p1;
13      if ( z )
14      {
15          if ( d >= 0 )  { S4(1,0,0,0);      Y_rec(p1+4, p2,  z); }   // 1 0 0 0
16          if ( d >= 2 )  { S6(1,0,1,0,0,0);  Y_rec(p1+6, p2, !z); }   // 1 0 1 0 0 0
17          if ( d >= 2 )  { S6(1,1,1,0,0,0);  Y_rec(p1+6, p2,  z); }   // 1 1 1 0 0 0
18          if ( d >= 1 )  { S5(1,1,0,0,0);    Y_rec(p1+5, p2, !z); }   // 1 1 0 0 0
19          if ( d >= 0 )  { S1(0);            Y_rec(p1+1, p2,  z); }   // 0
20      }
21      else
22      {
23          if ( d >= 0 )  { S1(0);            Y_rec(p1+1, p2,  z); }   // 0
24          if ( d >= 1 )  { S5(1,1,0,0,0);    Y_rec(p1+5, p2, !z); }   // 1 1 0 0 0
25          if ( d >= 2 )  { S6(1,1,1,0,0,0);  Y_rec(p1+6, p2,  z); }   // 1 1 1 0 0 0
26          if ( d >= 2 )  { S6(1,0,1,0,0,0);  Y_rec(p1+6, p2, !z); }   // 1 0 1 0 0 0
27          if ( d >= 0 )  { S4(1,0,0,0);      Y_rec(p1+4, p2,  z); }   // 1 0 0 0
28      }
29  }
```

Note the conditions `if ( d>= ? )` that make sure that no string appears repeated. The initial call is `Y_rec(0, n-1, 0)`. The sequence of the numbers $y(n)$ of length-n strings starts as

```
n:     0 1 2 3  4  5  6  7  8   9  10  11  12  13   14   15   16   17
y(n):  1 2 4 8 12 17 25 41 69 114 180 280 440 705 1137 1825 2905 4610
```

The generating function is

$$\sum_{n=0}^{\infty} y(n)\,x^n \quad = \quad \frac{1 + x + 2\,x^2 + 4\,x^3 + 3\,x^4 + 2\,x^5}{1 - x - x^4 - x^5 - 2\,x^6} \tag{12.8-5}$$

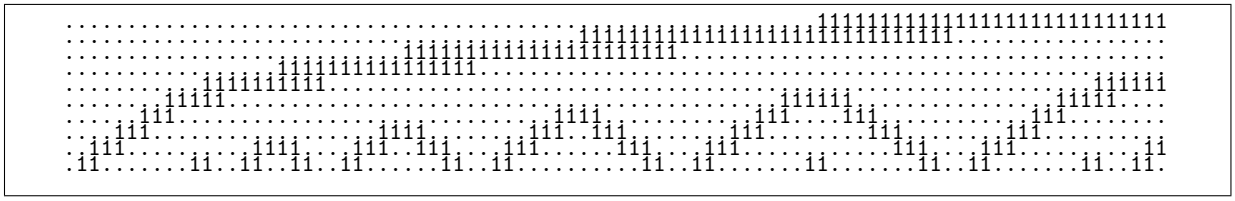

Figure 12.8-C: A Gray code for the length-10 binary strings with no substring 1x1 or 1xy1.

12.8.3 Neither substrings 1x1 nor substrings 1xy1

A recursion for a Gray code of the of n-bit binary words $Z(n)$ with no substrings 1x1 or 1xy1 (shown in figure 12.8-C) is

$$Z(n) \quad = \quad \begin{matrix} [1\,0\,0\,0\,.\,Z(n-4) \quad\quad\;\;] \\ [1\,1\,0\,0\,0\,.\,Z^{\mathbf{R}}(n-5)] \\ [0\,.\,Z(n-1) \quad\quad\;\;] \end{matrix} \tag{12.8-6}$$

The sequence of the numbers $z(n)$ of length-n strings starts as

```
n:     0 1 2 3 4  5  6  7  8  9 10  11  12  13  14  15   16   17
z(n):  1 2 4 6 8 11 17 27 41 60 88 132 200 301 449 669 1001 1502
```

The sequence is (apart from three leading ones) entry A079972 of [240] where two combinatorial interpretations are given:

```
Number of permutations satisfying -k<=p(i)-i<=r and p(i)-i not in I, i=1..n,
    with k=1, r=4, I={1,2}.
Number of compositions (ordered partitions) of n into elements of the set {1,4,5}.
```

The generating function is

$$\sum_{n=0}^{\infty} z(n)\,x^n \quad = \quad \frac{1+x+2\,x^2+2\,x^3+x^4}{1-x-x^4-x^5} \tag{12.8-7}$$

Chapter 13

Parenthesis strings

13.1 Co-lexicographic order

```
 1:  ((((()))))  11111.....      22:  (())(()())  11..11.1..
 2:  (((()(.)))  1111.1....      23:  ()()(()())  1.1.11.1..
 3:  (((()((.)))  111.11....     24:  ((()))(())  111...11..
 4:  (()(((.)))  11.111....      25:  (()())(())  11.1..11..
 5:  ()(((.)))  1.1111....       26:  ()(())(())  1.11..11..
 6:  ((((.))))  1111..1...       27:  (())()(())  11..1.11..
 7:  (((()()))  111.1.1...       28:  ()()()(())  1.1.1.11..
 8:  (()(()))  11.11.1...        29:  (((()))))()  1111....1.
 9:  ()(((.)))  1.111.1...       30:  (((.)))()  111.1...1.
10:  (((.)(.)))  111..11...      31:  (()(.))()  11.11...1.
11:  (()()(.)))  11.1.11...      32:  ()(((.)))()  1.111...1.
12:  ()(()(.)))  1.11.11...      33:  ((())())()  111..1..1.
13:  (())((.)))  11..111...      34:  (()()())()  11.1.1..1.
14:  ()()((.)))  1.1.111...      35:  ()(()())()  1.11.1..1.
15:  ((((.))))()  1111...1..     36:  (())(())()  11..11..1.
16:  (((.)(.)))()  111.1..1..    37:  ()()(())()  1.1.11..1.
17:  (()((.)))()  11.11..1..     38:  ((()))()()  111...1.1.
18:  ()(((.)))()  1.111..1..     39:  (()())()()  11.1..1.1.
19:  (((.))()())  111..1.1..     40:  ()(())()()  1.11..1.1.
20:  (()()(.)))()  11.1.1.1..    41:  (())()()()  11..1.1.1.
21:  ()(()(.)))()  1.11.1.1..    42:  ()()()()()  1.1.1.1.1.
```

Figure 13.1-A: All (42) valid strings of 5 pairs of parenthesis in colex order.

An iterative scheme to generate all valid ways to group parenthesis can be obtained from a modified version of the combinations in co-lexicographic order (see section 6.1.2 on page 177). For $n = 5$ pairs the possible combinations are shown in figure 13.1-A. This is the output of [FXT: comb/paren-demo.cc].

Consider the sequences to the right of the paren strings as binary words. Whenever the leftmost block has more than one bit then its rightmost bit is moved one position to the right. If the leftmost block consists of a single bit then the bits of the longest run of the repeated pattern '`1.`' at the left are gathered at the left end and further, the rightmost bit in next block of ones (which contains at least two ones) is moved by one position to the right and the rest of the block is gathered at the left end (see the transitions from `#13` to `#14` or `#36` to `#37`).

A sentinel `x[k]` is used to save one branch with the generation of the next string [FXT: `class paren` in comb/paren.h]:

```
1   class paren
2   {
3   public:
4       ulong k_;  // Number of paren pairs
5       ulong n_;  // ==2*k
```

```
6       ulong *x_; // Positions where a opening paren occurs
7       char *str_;  // String representation,  e.g. "((())())()"
8
9   public:
10      paren(ulong k)
11      {
12          k_ = (k>1 ? k : 2);  // not zero (empty) or one (trivial: "()")
13          n_ = 2 * k_;
14          x_ = new ulong[k_ + 1];
15          x_[k_] = 999;  // sentinel
16          str_ = new char[n_ + 1];
17          str_[n_] = 0;
18          first();
19      }
20
21      ~paren()
22      {
23          delete [] x_;
24          delete [] str_;
25      }
26
27      void first()  { for (ulong i=0; i<k_; ++i)  x_[i] = i; }
28
29      void last()  { for (ulong i=0; i<k_; ++i)  x_[i] = 2*i; }
30
31      [--snip--]
```

The code for the computation of the successor and predecessor is quite concise:

```
1       ulong next()  // return zero if current paren is the last
2       {
3           // if ( k_==1 )  return 0;  // uncomment to make algorithm work for k_==1
4
5           ulong j = 0;
6           if ( x_[1] == 2 )
7           {
8               // scan for low end == 010101:
9               j = 2;
10              while ( (j<=k_) && (x_[j]==2*j) )  ++j;  // can touch sentinel
11              if ( j==k_ )  {  first();  return 0; }
12          }
13
14          // scan block:
15          while ( 1 == (x_[j+1] - x_[j]) )  { ++j; }
16
17          ++x_[j];  // move edge element up
18          for (ulong i=0; i<j; ++i)  x_[i] = i; // attach block at low end
19
20          return 1;
21      }
22
23      ulong prev()  // return zero if current paren is the first
24      {
25          // if ( k_==1 )  return 0;  // uncomment to make algorithm work for k_==1
26
27          ulong j = 0;
28          // scan for first gap:
29          while ( x_[j]==j )  ++j;
30          if ( j==k_ )  {  last();  return 0; }
31
32          if ( x_[j]-x_[j-1] == 2 )  --x_[j];  // gap of length one
33          else
34          {
35              ulong i = --x_[j];
36              --j;
37              --i;
38              // j items to go, distribute as 1.1.1.11111
39              for (  ;  2*i>j;  --i,--j)  x_[j] = i;
40              for (  ; i; --i)  x_[i] = 2*i;
41              x_[0] = 0;
42          }
43
44          return 1;
45      }
46
47      const ulong * data()  { return x_; }
48      [--snip--]
49
```

The strings are set up on demand only:

```
1       const char * string()  // generate on demand
2       {
3           for (ulong j=0; j<n_; ++j)  str_[j] = ')';
4           for (ulong j=0; j<k_; ++j)  str_[x_[j]] = '(';
5           return str_;
6       }
7   };
```

The 477,638,700 paren words for $n = 18$ are generated at a rate of about 65 million objects per second. Section 1.31 on page 81 gives a bit-level algorithm for the generation of the paren words in colex order.

13.2 Gray code via restricted growth strings

```
 1:    [ 0, 0, 0, 0, ]    ()()()()    1.1.1.1.
 2:    [ 0, 0, 0, 1, ]    ()()(())    1.1.11..
 3:    [ 0, 0, 1, 0, ]    ()(())()    1.11..1.
 4:    [ 0, 0, 1, 1, ]    ()(()())    1.11.1..
 5:    [ 0, 0, 1, 2, ]    ()((()))    1.111...
 6:    [ 0, 1, 0, 0, ]    (())()()    11..1.1.
 7:    [ 0, 1, 0, 1, ]    (())(())    11..11..
 8:    [ 0, 1, 1, 0, ]    (()())()    11.1..1.
 9:    [ 0, 1, 1, 1, ]    (()()())    11.1.1..
10:    [ 0, 1, 1, 2, ]    (()(()))    11.11...
11:    [ 0, 1, 2, 0, ]    ((()))()    111...1.
12:    [ 0, 1, 2, 1, ]    ((())())    111..1..
13:    [ 0, 1, 2, 2, ]    (((())))    111.1...
14:    [ 0, 1, 2, 3, ]    ((((0)))    1111....
```

Figure 13.2-A: Length-4 restricted growth strings (left), and the corresponding paren strings (middle) and delta sets (right).

The valid paren strings can be represented by sequences $a_0, a_1, \ldots, a_n$ where $a_0 = 0$ and $a_k \leq a_{k-1} + 1$. These sequences are examples of so-called *restricted growth strings* (RGS). Some sources use the term *restricted growth functions*. The RGSs for $n = 4$ are shown in figure 13.2-A. A RGS can be incremented by incrementing the highest (rightmost in figure 13.2-A) digit a_j where $a_j \leq a_{j-1}$ and setting $a_i = 0$ for all $i > j$. A decrement is obtained by decrementing the highest digit $a_j \neq 0$ and setting $a_i = a_{i-1} + 1$ for all $i > j$.

The RGSs for a given n can be generated as follows [FXT: `class catalan` in comb/catalan.h]:

```
1   class catalan
2   // Catalan restricted growth strings (RGS)
3   // By default in near-perfect minimal-change order, i.e.
4   //  exactly two symbols in paren string change with each step
5   {
6   public:
7       int *as_;   // digits of the RGS: as_[k] <= as[k-1] + 1
8       int *d_;    // direction with recursion (+1 or -1)
9       ulong n_;   // Number of digits (paren pairs)
10      char *str_; // paren string
11      bool xdr_;  // whether to change direction in recursion (==> minimal-change order)
12      int dr0_;   // dr0: starting direction in each recursive step:
13      //   dr0=+1  ==> start with as[]=[0,0,0,...,0]   == "()()()...()"
14      //   dr0=-1  ==> start with as[]=[0,1,2,...,n-1] == "((( ... )))"
```

```
1   public:
2       catalan(ulong n, bool xdr=true, int dr0=+1)
3           : n_(n)
4       {
5           as_ = new int[n_];
6           d_ = new int[n_];
7           str_ = new char[2*n_+1];  str_[2*n_] = 0;
8           init(xdr, dr0);
```

```
     1:  [ 0 1 2 3 4 5 ]    [ - - - - - - ]    ((((()))))    11111.....
     2:  [ 0 1 2 3 3 3 ]    [ - - - - - - ]    ((((()))))    1111.1....    ((((XA))))
     3:  [ 0 1 2 3 3 2 ]    [ - - - - - - ]    ((((())))     1111..1...    ((((()XA)))
     4:  [ 0 1 2 3 3 1 ]    [ - - - - - - ]    ((((()))()    1111...1..    ((((())XA))
     5:  [ 0 1 2 3 3 0 ]    [ - - - - - - ]    ((((())))()   1111....1.    ((((()))XA)
     6:  [ 0 1 2 2 2 0 ]    [ - - - - - + ]    (((()))))()   111.1...1.    (((XA)))()
     7:  [ 0 1 2 2 2 1 ]    [ - - - - - + ]    (((()()())))   111.1..1..    (((()())AX)
     8:  [ 0 1 2 2 2 2 ]    [ - - - - - + ]    (((()()()))    111.1.1...    (((()()AX))
     9:  [ 0 1 2 2 2 3 ]    [ - - - - - + ]    (((()))))      111.111....    (((()(AX)))
    10:  [ 0 1 2 1 2 2 ]    [ - - - - - - ]    (((())(())     111..11...    (((()X(A)))   2
    11:  [ 0 1 2 1 1 1 ]    [ - - - - - - ]    (((())()()     111..1.1..    ((((())(XA))
    12:  [ 0 1 2 1 1 0 ]    [ - - - - - - ]    (((())())()    111..1..1.    ((((())()XA)
    13:  [ 0 1 2 0 0 0 ]    [ - - - - - + ]    ((((()))()()   111...1.1.    ((((()XA)()
    14:  [ 0 1 2 0 1 1 ]    [ - - - - - + ]    ((((()))(()    111...11..    ((((()))(AX)
    15:  [ 0 1 1 0 1 2 ]    [ - - - + - + ]    (()(()))(())   11.1..11..    ((XA))(())
    16:  [ 0 1 1 0 0 0 ]    [ - - - + - + ]    (()(()))()()   11.1..1.1.    (()())(XA)
    17:  [ 0 1 1 1 1 0 ]    [ - - - + + + ]    (()()())()()   11.1.1..1.    (()()AX)()
    18:  [ 0 1 1 1 1 1 ]    [ - - - + + + ]    (()()()))()    11.1.1.1..    (()()()AX)
    19:  [ 0 1 1 1 2 2 ]    [ - - - + + + ]    (()()(()))     11.1.11...    (()()(AX))
    20:  [ 0 1 1 2 3 3 ]    [ - - - + - + ]    (()((())))     11.111....    (()(A(X)))   2
    21:  [ 0 1 1 2 2 2 ]    [ - - - + - + ]    (()(()()))     11.11.1...    (()((XA)))
    22:  [ 0 1 1 2 1 1 ]    [ - - - + - + ]    (()(())())     11.11..1..    (()(()XA))
    23:  [ 0 1 1 2 0 0 ]    [ - - - + - + ]    (()(()))()     11.11...1.    (()(())XA)
    24:  [ 0 1 0 1 0 1 ]    [ - - - - - + ]    (())(())()     11..11..1.    (()X(A))()   2
    25:  [ 0 1 0 1 1 1 ]    [ - - - - - + ]    (())(()()))    11..11.1..    (())(()AX)
    26:  [ 0 1 0 1 1 2 ]    [ - - - - - + ]    (())(((())))   11..111...    (())((AX))
    27:  [ 0 1 0 0 1 1 ]    [ - - - - - - ]    (())()(())     11..1.11..    (())(X(A))   2
    28:  [ 0 1 0 0 0 0 ]    [ - - - - - - ]    (())()()()     11..1.1.1.    (())()(XA)
    29:  [ 0 0 0 0 0 0 ]    [ - - + + + + ]    ()()()()()     1.1.1.1.1.    (XA)()()()
    30:  [ 0 0 0 0 0 1 ]    [ - - + + + + ]    ()()()(()))    1.1.1.11..    ()()()(AX)
    31:  [ 0 0 0 0 1 2 ]    [ - - + + + - ]    ()()((())))    1.1.111...    ()()(A(X))   2
    32:  [ 0 0 0 0 1 1 ]    [ - - + + + - ]    ()()(()()))    1.1.11.1..    ()()((XA))
    33:  [ 0 0 0 0 1 0 ]    [ - - + + + - ]    ()()(())()     1.1.11..1.    ()()(()XA)
    34:  [ 0 0 1 2 2 0 ]    [ - - + + - + ]    ()((()))()     1.111...1.    ()(A(X))()   2
    35:  [ 0 0 1 2 2 1 ]    [ - - + + - + ]    ()((()))())    1.111..1..    ()((())AX)
    36:  [ 0 0 1 2 2 2 ]    [ - - + + - + ]    ()((()()))     1.111.1...    ()((()AX))
    37:  [ 0 0 1 2 3 3 ]    [ - - + + - + ]    ()(((())))     1.1111....    ()(((AX)))
    38:  [ 0 0 1 1 1 2 ]    [ - - + + - - ]    ()(()(()))     1.11.11...    ()((X(A)))   2
    39:  [ 0 0 1 1 1 1 ]    [ - - + + - - ]    ()(()()()))    1.11.1.1..    ()(()(XA))
    40:  [ 0 0 1 1 1 0 ]    [ - - + + - - ]    ()(()())()     1.11.1..1.    ()(()()XA)
    41:  [ 0 0 1 0 0 0 ]    [ - - + + - + ]    ()(())()()     1.11..1.1.    ()(())XA)()
    42:  [ 0 0 1 0 1 1 ]    [ - - + + - + ]    ()(())(())     1.11..11..    ()(())(AX)
```

Figure 13.2-B: Minimal-change order for the paren strings of 5 pairs. From left to right: restricted growth strings, arrays of directions, paren strings, delta sets, and difference strings. If the change is not adjacent, then the distance of changed positions is given at the right. The order corresponds to `dr0=-1`.

```
 1:  [ 0 0 0 0 0 ]   [ + + + + + ]   ()()()()()   1.1.1.1.1.
 2:  [ 0 0 0 0 1 ]   [ + + + + + ]   ()()()(())   1.1.1.11..   ()()()(AX)
 3:  [ 0 0 0 1 2 ]   [ + + + + - ]   ()()((()))   1.1.111...   ()()(A(X))  2
 4:  [ 0 0 0 1 1 ]   [ + + + + - ]   ()()(()())   1.1.11.1..   ()()((XA))
 5:  [ 0 0 0 1 0 ]   [ + + + + - ]   ()()(())()   1.1.11..1.   ()()(()XA)
 6:  [ 0 0 1 2 0 ]   [ + + + - + ]   ()((()))()   1.111...1.   ()(A(X))()  2
 7:  [ 0 0 1 2 1 ]   [ + + + - + ]   ()((())())   1.111..1..   ()((())AX)
 8:  [ 0 0 1 2 2 ]   [ + + + - + ]   ()((()()))   1.111.1...   ()((()AX))
 9:  [ 0 0 1 2 3 ]   [ + + + - + ]   ()(((())))   1.1111....   ()(((AX)))
10:  [ 0 0 1 1 2 ]   [ + + + - - ]   ()(()(()))   1.11.11...   ()((X(A)))  2
11:  [ 0 0 1 1 1 ]   [ + + + - - ]   ()(()()())   1.11.1.1..   ()(()(XA))
12:  [ 0 0 1 1 0 ]   [ + + + - - ]   ()(()())()   1.11.1..1.   ()(()()XA)
13:  [ 0 0 1 0 0 ]   [ + + + - + ]   ()(())()()   1.11..1.1.   ()(()XA)()
14:  [ 0 0 1 0 1 ]   [ + + + - + ]   ()(())(())   1.11..11..   ()(())(AX)
15:  [ 0 1 2 0 1 ]   [ + + - + - ]   ((()))(())   111...11..   (A(X))(())  2
16:  [ 0 1 2 0 0 ]   [ + + - + - ]   ((()))()()   111...1.1.   ((()))(XA)
17:  [ 0 1 2 1 0 ]   [ + + - + + ]   ((()))(())   111..1..1.   ((()))AX)()
18:  [ 0 1 2 1 1 ]   [ + + - + + ]   ((()))()()   111..1.1..   ((()))()AX)
19:  [ 0 1 2 1 2 ]   [ + + - + + ]   ((()))(())   111..11...   ((()))(AX))
20:  [ 0 1 2 2 3 ]   [ + + - + - ]   (((()(())))   111.11....   (((()A(X)))  2
21:  [ 0 1 2 2 2 ]   [ + + - + - ]   (((()()())))   111.1.1...   (((()(XA)))
22:  [ 0 1 2 2 1 ]   [ + + - + - ]   (((()())()))   111.1..1..   (((()()XA))
23:  [ 0 1 2 2 0 ]   [ + + - + - ]   (((())))()   111.1...1.   (((())()XA)
24:  [ 0 1 2 3 0 ]   [ + + - + + ]   ((((()))))()   1111....1.   (((AX)))()
25:  [ 0 1 2 3 1 ]   [ + + - + + ]   ((((()))()))   1111...1..   ((((()))AX)
26:  [ 0 1 2 3 2 ]   [ + + - + + ]   ((((())(())))   1111..1...   ((((())AX))
27:  [ 0 1 2 3 3 ]   [ + + - + + ]   ((((()))))   1111.1....   ((((AX)))
28:  [ 0 1 2 3 4 ]   [ + + - + + ]   (((((0)))))   11111.....   ((((AX))))
29:  [ 0 1 1 2 3 ]   [ + + - - - ]   (()(((0)))   11.111....   ((X((A))))  3
30:  [ 0 1 1 2 2 ]   [ + + - - - ]   (()(()()))   11.11.1...   (()((XA)))
31:  [ 0 1 1 2 1 ]   [ + + - - - ]   (()(())())   11.11..1..   (()(()XA))
32:  [ 0 1 1 2 0 ]   [ + + - - - ]   (()(()))()   11.11...1.   (()(())XA)
33:  [ 0 1 1 1 0 ]   [ + + - - + ]   (()()())()   11.1.1..1.   (()(XA))()
34:  [ 0 1 1 1 1 ]   [ + + - - + ]   (()()()())   11.1.1.1..   (()()(AX)
35:  [ 0 1 1 1 2 ]   [ + + - - + ]   (()()(()))   11.1.11...   (()()(AX))
36:  [ 0 1 1 0 1 ]   [ + + - - - ]   (()())(())   11.1..11..   (()()X(A))  2
37:  [ 0 1 1 0 0 ]   [ + + - - - ]   (()())()()   11.1..1.1.   (()())(XA)
38:  [ 0 1 0 0 0 ]   [ + + - + + ]   (())()()()   11..1.1.1.   (()XA)()()
39:  [ 0 1 0 0 1 ]   [ + + - + + ]   (())()(())   11..1.11..   (())()(AX)
40:  [ 0 1 0 1 2 ]   [ + + - + - ]   (())((()))   11..111...   (())(A(X))  2
41:  [ 0 1 0 1 1 ]   [ + + - + - ]   (())(()())   11..11.1..   (())((XA))
42:  [ 0 1 0 1 0 ]   [ + + - + - ]   (())(())()   11..11..1.   (())(()XA)
```

Figure 13.2-C: Minimal-change order for the paren strings of 5 pairs. From left to right: restricted growth strings, arrays of directions, paren strings, delta sets, and difference strings. If the change is not adjacent, then the distance of changed positions is given at the right. The order corresponds to `dr0=+1`.

```
9       }
10
11      ~catalan()
12      {
13          delete [] as_;
14          delete [] d_;
15          delete [] str_;
16      }
17
18      void init(bool xdr, int dr0)
19      {
20          dr0_ = ( (dr0>=0) ? +1 : -1 );
21          xdr_ = xdr;
22
23          ulong n = n_;
24          if ( dr0_>0 )  for (ulong k=0; k<n; ++k)  as_[k] = 0;
25          else           for (ulong k=0; k<n; ++k)  as_[k] = k;
26
27          for (ulong k=0; k<n; ++k)  d_[k] = dr0_;
28      }
29
30      bool next()  { return next_rec(n_-1); }
31
32      const int *get()  const  { return as_; }
33
34      const char* str()  { make_str();  return (const char*)str_; }
35
36     [--snip--]
```

The minimal-change order is obtained by changing the 'direction' in the recursion, an essentially identical mechanism (for the generation of set partitions) is shown in chapter 15 on page 335. The function is given in [FXT: comb/catalan.cc]:

```
1   bool
2   catalan::next_rec(ulong k)
3   {
4       if ( k<1 )  return false;  // current is last
5
6       int d = d_[k];
7       int as = as_[k] + d;
8       bool ovq = ( (d>0) ? (as>as_[k-1]+1) : (as<0) );
9       if ( ovq )  // have to recurse
10      {
11          ulong ns1 = next_rec(k-1);
12          if ( 0==ns1 )  return false;
13
14          d = ( xdr_ ? -d : dr0_ );
15          d_[k] = d;
16
17          as = ( (d>0) ? 0 : as_[k-1]+1 );
18      }
19      as_[k] = as;
20
21      return true;
22  }
```

The program [FXT: comb/catalan-demo.cc] demonstrates the usage:

```
ulong n = 4;
bool xdr = true;
int dr0 = -1;

catalan C(n, xdr, dr0);
do
{
    // visit
}
while ( C.next() );
```

About 67 million strings per second are generated. Figure 13.2-B shows the minimal-change for $n = 5$ and `dr0=-1`, figure 13.2-C shows the order for `dr0=+1`.

More minimal-change orders

The Gray code order shown in figure 13.2-D can be generated via a simple recursion:

```
 1:   0 0 0 0 0    1.1.1.1.1.        22:   0 1 2 2 1    111.1..1..
 2:   0 0 0 0 1    1.1.1.11..        23:   0 1 2 2 0    111.1...1.
 3:   0 0 0 1 2    1.1.111...        24:   0 1 2 1 0    111..1..1.
 4:   0 0 0 1 1    1.1.11.1..        25:   0 1 2 1 1    111..1.1..
 5:   0 0 0 1 0    1.1.11..1.        26:   0 1 2 1 2    111..11...
 6:   0 0 1 2 3    1.1111....        27:   0 1 2 0 1    111...11..
 7:   0 0 1 2 2    1.111.1...        28:   0 1 2 0 0    111...1.1.
 8:   0 0 1 2 1    1.111..1..        29:   0 1 1 0 0    11.1..1.1.
 9:   0 0 1 2 0    1.111...1.        30:   0 1 1 0 1    11.1..11..
10:   0 0 1 1 0    1.11.1..1.        31:   0 1 1 1 2    11.1.11...
11:   0 0 1 1 1    1.11.1.1..        32:   0 1 1 1 1    11.1.1.1..
12:   0 0 1 1 2    1.11.11...        33:   0 1 1 1 0    11.1.1..1.
13:   0 0 1 0 1    1.11..11..        34:   0 1 1 2 0    11.11...1.
14:   0 0 1 0 0    1.11..1.1.        35:   0 1 1 2 1    11.11..1..
15:   0 1 2 3 0    1111....1.        36:   0 1 1 2 2    11.11.1...
16:   0 1 2 3 1    1111...1..        37:   0 1 1 2 3    11.111....
17:   0 1 2 3 2    1111..1...        38:   0 1 0 1 0    11..11..1.
18:   0 1 2 3 3    1111.1....        39:   0 1 0 1 1    11..11.1..
19:   0 1 2 3 4    11111.....        40:   0 1 0 1 2    11..111...
20:   0 1 2 2 3    111.11....        41:   0 1 0 0 1    11..1.11..
21:   0 1 2 2 2    111.1.1...        42:   0 1 0 0 0    11..1.1.1.
```

Figure 13.2-D: Strings of 5 pairs of parenthesis in a Gray code order as generated by a recursive algorithm.

```
1   ulong n;    // Number of paren pairs
2   ulong *rv;  // restricted growth strings
3
4   void next_rec(ulong d, bool z)
5   {
6       if ( d==n )  visit();
7       else
8       {
9           const long rv1 = rv[d-1];  // left neighbor
10          if ( 0==z )
11          {
12              for (long x=0; x<=rv1+1; ++x)  // forward
13              {
14                  rv[d] = x;
15                  next_rec(d+1, (x&1));
16              }
17          }
18          else
19          {
20              for (long x=rv1+1; x>=0; --x)  // backward
21              {
22                  rv[d] = x;
23                  next_rec(d+1, !(x&1));
24              }
25          }
26      }
27  }
```

The initial call is `next_rec(0, 0);`. About 81 million strings per second are generated [FXT: comb/paren-gray-rec-demo.cc].

An loopless algorithm (that does not use RGS) given in [251] is implemented in [FXT: `class paren_gray` in comb/paren-gray.h]. The generated order for five paren pairs is shown in figure 13.2-E About 54 million strings per second are generated [FXT: comb/paren-gray-demo.cc]. Still more algorithms for the parentheses strings in minimal-change order are given in [74], [258], and [278].

For even values of n it is possible to generate paren strings in *strong minimal-change order* where changes occur only in adjacent positions. Figure 13.2-F shows an example for four pairs of parens. The listing was generated with [FXT: graph/graph-parengray-demo.cc] that uses directed graphs and the search algorithms described in chapter 19 on page 373.

```
 1:  ()()()()()    1.1.1.1.1.      22:  (()()(()))    11.1.11...
 2:  ()()()(())    1.1.1.11..      23:  (()()())()    11.1.1..1.
 3:  ()()(())()    1.1.11.1..      24:  (())()()()    111..1..1.
 4:  ()()((()))    1.1.111...      25:  (())()(())    111..11...
 5:  ()()(())()    1.1.11..1.      26:  (())(())()    111..1.1..
 6:  ()(())()()    1.11.1..1.      27:  (())(()())    111.1..1..
 7:  ()(())(())    1.11.11...      28:  (())(())()    111.1.1...
 8:  ()(()()())    1.11.1.1..      29:  (()()(()))    111.11....
 9:  ()((()))()    1.111..1..      30:  (()()())()    111.1...1.
10:  ()((()()))    1.111.1...      31:  (()(()))()    1111....1.
11:  ()(((())))    1.1111....      32:  ((((()))))    11111.....
12:  ()((()))()    1.111...1.      33:  ((()()))()    1111.1....
13:  ()(())()()    1.11..1.1.      34:  ((())())()    1111..1...
14:  ()(())(())    1.11..11..      35:  (((()))())    1111...1..
15:  (())()(())    11.1..11..      36:  ((()))(())    111...11..
16:  (())()()()    11.1..1.1.      37:  ((()))()()    111...1.1.
17:  (())(())()    11.11...1.      38:  (())()()()    11..1.1.1.
18:  (())((()))    11.111....      39:  (())()(())    11..1.11..
19:  (())(()())    11.11.1...      40:  (())(())()    11..11.1..
20:  (()(()))()    11.11..1..      41:  (()))(()))    11..111...
21:  (()()()())    11.1.1.1..      42:  (())(())()    11..11..1.
```

Figure 13.2-E: Strings of 5 pairs of parenthesis in Gray code order as generated by a loopless algorithm.

```
 0:   ....1111 == (((())))
 1:   ...1.111 == (((()(   ^= ...11...
 2:   ...11.11 == (()(())   ^= ....11..
 3:   ...111.1 == ()((()))   ^= .....11.
 4:   ..1.11.1 == ()(()())   ^= ..11....
 5:   ..1.1.11 == (()()())   ^= .....11.
 6:   ..1..111 == (((),))   ^= ....11..
 7:   .1...111 == ((()))()   ^= .11.....
 8:   .1..1.11 == (()())()   ^= ....11..
 9:   .1..11.1 == ()(())()   ^= .....11.
10:   .1.1.1.1 == ()()()()   ^= ...11...
11:   ..11.1.1 == ()()(())   ^= .11.....
12:   ..11..11 == (())(())   ^= .....11.
13:   .1.1..11 == (())()()   ^= .11.....
```

Figure 13.2-F: A strong minimal-change order for the paren strings of 4 pairs.

13.3 The number of parenthesis strings: Catalan numbers

The number of valid combinations of n parenthesis pairs is

$$C_n \;=\; \frac{\binom{2n}{n}}{n+1} \;=\; \frac{\binom{2n+1}{n}}{2n+1} \;=\; \frac{\binom{2n-2}{n-1}}{n} \;=\; \binom{2n}{n} - \binom{2n}{n-1} \qquad (13.3\text{-}1)$$

as nicely explained in [133, p.343-346]. These are the *Catalan numbers*, sequence A000108 of [240]:

n :	C_n	n :	C_n	n :	C_n
1:	1	11:	58786	21:	24466267020
2:	2	12:	208012	22:	91482563640
3:	5	13:	742900	23:	343059613650
4:	14	14:	2674440	24:	1289904147324
5:	42	15:	9694845	25:	4861946401452
6:	132	16:	35357670	26:	18367353072152
7:	429	17:	129644790	27:	69533550916004
8:	1430	18:	477638700	28:	263747951750360
9:	4862	19:	1767263190	29:	1002242216651368
10:	16796	20:	6564120420	30:	3814986502092304

The Catalan numbers are generated most easily with the relation

$$C_{n+1} \;=\; \frac{2\,(2\,n+1)}{n+2}\,C_n \tag{13.3-2}$$

The generating function is

$$C(x) \;=\; \frac{1-\sqrt{1-4\,x}}{2\,x} \;=\; \sum_{n=0}^{\infty} C_n\,x^n \;=\; 1+x+2\,x^2+5\,x^3+14\,x^4+42\,x^5+\ldots \tag{13.3-3}$$

One further has the convolution property $[x\,C(x)] = x + [x\,C(x)]^2$

$$C_n \;=\; \sum_{k=0}^{n-1} C_k\,C_{n-1-k} \tag{13.3-4}$$

13.4 Increment-i RGS and k-ary trees

13.4.1 Generation in lexicographic order

```
 1:  [ 0 0 0 0 ]    21:  [ 0 1 2 1 ]    41:  [ 0 2 2 3 ]
 2:  [ 0 0 0 1 ]    22:  [ 0 1 2 2 ]    42:  [ 0 2 2 4 ]
 3:  [ 0 0 0 2 ]    23:  [ 0 1 2 3 ]    43:  [ 0 2 3 0 ]
 4:  [ 0 0 1 0 ]    24:  [ 0 1 2 4 ]    44:  [ 0 2 3 1 ]
 5:  [ 0 0 1 1 ]    25:  [ 0 1 3 0 ]    45:  [ 0 2 3 2 ]
 6:  [ 0 0 1 2 ]    26:  [ 0 1 3 1 ]    46:  [ 0 2 3 3 ]
 7:  [ 0 0 1 3 ]    27:  [ 0 1 3 2 ]    47:  [ 0 2 3 4 ]
 8:  [ 0 0 2 0 ]    28:  [ 0 1 3 3 ]    48:  [ 0 2 3 5 ]
 9:  [ 0 0 2 1 ]    29:  [ 0 1 3 4 ]    49:  [ 0 2 4 0 ]
10:  [ 0 0 2 2 ]    30:  [ 0 1 3 5 ]    50:  [ 0 2 4 1 ]
11:  [ 0 0 2 3 ]    31:  [ 0 2 0 0 ]    51:  [ 0 2 4 2 ]
12:  [ 0 0 2 4 ]    32:  [ 0 2 0 1 ]    52:  [ 0 2 4 3 ]
13:  [ 0 1 0 0 ]    33:  [ 0 2 0 2 ]    53:  [ 0 2 4 4 ]
14:  [ 0 1 0 1 ]    34:  [ 0 2 1 0 ]    54:  [ 0 2 4 5 ]
15:  [ 0 1 0 2 ]    35:  [ 0 2 1 1 ]    55:  [ 0 2 4 6 ]
16:  [ 0 1 1 0 ]    36:  [ 0 2 1 2 ]
17:  [ 0 1 1 1 ]    37:  [ 0 2 1 3 ]
18:  [ 0 1 1 2 ]    38:  [ 0 2 2 0 ]
19:  [ 0 1 1 3 ]    39:  [ 0 2 2 1 ]
20:  [ 0 1 2 0 ]    40:  [ 0 2 2 2 ]
```

Figure 13.4-A: The 55 increment-2 restricted growths strings of length 4.

We now allow an increment of i in the restricted growth strings ($i = 1$ corresponds to the paren RGS of section 13.2). Figure 13.4-A shows the increment-2 restricted growths strings of length 4. The strings can be generated in lexicographic order via [FXT: `class rgs_binomial` in comb/rgs-binomial.h].

```
1   class rgs_binomial
2   // Restricted growth strings (RGS) s[0,...,n-1] so that s[k] <= s[k-1]+i
3   {
4   public:
5       ulong *s_;  // restricted growth string
6       ulong n_;   // Length of strings
7       ulong i_;   // s[k] <= s[k-1]+i
8     [--snip--]
9
10      ulong next()
11      // Return index of first changed element in s[],
12      // Return zero if current string is the last
13      {
14          ulong k = n_;
15
16      start:
17          --k;
```

```
18          if ( k==0 )  return 0;
19
20          ulong sk = s_[k] + 1;
21          ulong mp = s_[k-1] + i_;
22          if ( sk > mp )  // "carry"
23          {
24              s_[k] = 0;
25              goto start;
26          }
27
28          s_[k] = sk;
29          return k;
30      }
31    [--snip--]
```

The rate of generation is about 129 M/s for $i = 1$ (corresponding to paren strings), 143 M/s for $i = 2$, and 156 M/s with $i = 3$ [FXT: comb/rgs-binomial-demo.cc].

13.4.2 The number of increment-i RGS

The number $C_{n,i}$ of length-n increment-i strings equals

$$C_{n,i} \quad = \quad \frac{\binom{(i+1)\,n}{n}}{i\,n+1} \tag{13.4-1}$$

A recursion generalizing relation 13.3-2 is

$$C_{n+1,i} \quad = \quad (i+1)\,\frac{\prod_{k=1}^{i}\left[(i+1)\,n+k\right]}{\prod_{k=1}^{i}\left[i\,n+k+1\right]}\,C_{n,i} \tag{13.4-2}$$

The sequences of numbers of length-n strings for $i = 1, 2, 3, 4$ start

```
     n: 1 2  3   4    5     6      7       8        9         10          11
  i=1:  1 2  5  14   42   132    429    1430     4862     16796       58786
  i=2:  1 3 12  55  273  1428   7752   43263   246675   1430715     8414640
  i=3:  1 4 22 140  969  7084  53820  420732  3362260  27343888   225568798
  i=4:  1 5 35 285 2530 23751 231880 2330445 23950355 250543370 2658968130
```

These are respectively the entries A000108, A001764, A002293, A002294 of [240] where combinatorial interpretations are given. We note that for the generating function $C_i(x)$ we have the following expression as a hypergeometric function (see section 35.2 on page 685):

$$C_i(x) \quad = \quad \sum_{n=0}^{\infty} C_{n,i}\,x^n \tag{13.4-3a}$$

$$\quad = \quad F\left(\begin{matrix} 1/(i+1),\ 2/(i+1),\ 3/(i+1),\ \ldots,\ (i+1)/(i+1) \\ 2/i,\ 3/i,\ \ldots,\ i/i,\ (i+1)/i \end{matrix}\,\middle|\,\frac{(i+1)^{(i+1)}}{i^i}\,x\right) \tag{13.4-3b}$$

Note that the last upper and second last lower parameter cancel. Now let $f_i(x) := x\,C_i(x^i)$, then

$$f_i(x) - f_i(x)^{i+1} \quad = \quad x \tag{13.4-4}$$

That is, $f_i(x)$ can be obtained as the series reversion of $x^{i+1} - x$. We choose $i = 2$ for an example:

```
? t1=serreverse(x-x^3+O(x^(17)))
  x + x^3 + 3*x^5 + 12*x^7 + 55*x^9 + 273*x^11 + 1428*x^13 + 7752*x^15 + O(x^17)
? t2=hypergeom([1/3,2/3,3/3],[2/2,3/2],3^3/2^2*x)+O(x^17)
  1 + x + 3*x^2 + 12*x^3 + 55*x^4 + 273*x^5 + 1428*x^6 + 7752*x^7 + ... + O(x^17)
? f=x*subst(t2,x,x^2);
? t1-f
  O(x^17)      \\ f is actually the series reversion of x-x^3
? f-f^3
  x + O(x^35)  \\ ... so  f - f^3 == id
```

We further have the following convolution property, generalizing relation 13.3-4,

$$C_{n,i} \quad = \quad \sum_{j_1+j_2+\ldots+j_i+j_{(i+1)}=n-1} C_{j_1,i}\, C_{j_2,i}\, C_{j_3,i} \cdots C_{j_i,i}\, C_{j_{(i+1)},i} \tag{13.4-5}$$

An explicit expression for the function $C_i(x)$ is

$$C_i(x) \quad = \quad \exp\left(\frac{1}{i+1} \sum_{n=1}^{\infty} \binom{(i+1)\,n}{n} \frac{x^n}{n} \right) \tag{13.4-6}$$

The expression generalizes a relation given in [180, rel.6] (set $i = 1$ and take logarithm)

$$\sum_{n=1}^{\infty} \frac{1}{n} \binom{2\,n}{n} x^n \quad = \quad 2 \log\left(\frac{1-\sqrt{1-4\,x}}{2\,x} \right) \tag{13.4-7}$$

13.4.3 Gray code for k-ary trees

The length-n increment-i RGS correspond to k-ary trees with n internal nodes and $k = i+1$. An loopless algorithm for the generation of a Gray code for k-ary tress with only homogeneous changes is given in [29]. The RGS used in the algorithm gives the positions (one-based) of the ones in the delta sets, see figure 13.4-B. An implementation is [FXT: `class tree_gray` in comb/tree-gray.h]:

```
1   class tree_gray
2   {
3   public:
4       ulong *sq_;  // sequence of bit positions (seq[])  elements \in {1,2,...,n}
5       ulong *dr_;  // aux: direction            (dir[])
6       ulong *np_;  // aux: next position        (nextPos[])
7       ulong *mx_;  // aux: max position         (max[])
8       ulong n_;    // n (internal) nodes
9       ulong k_;    // k-ary tree
10
11      tree_gray(ulong n, ulong k)
12      {
13          n_ = n;
14          k_ = k;
15
16          // all arrays are one-based
17          sq_ = new ulong[n_+1];
18          dr_ = new ulong[n_+1];
19          np_ = new ulong[n_+2];  // one pad element right
20          mx_ = new ulong[n_+1];  // unchanged in next()
21          first();
22      }
23    [--snip--]
```

```
1       void first(ulong k=0)
2       {
3           if ( k )  k_ = k;
4           for (ulong j=1, e=1; j<=n_; ++j, e+=k_)  sq_[j] = mx_[j] = e;
5           for (ulong j=0; j<=n_; ++j)  dr_[j] = 1;  // "right"
6           for (ulong j=0; j<=n_+1; ++j)  np_[j] = j - 1;
7       }
```

The computation of the successor is a variant of the method given in [43]:

```
1       ulong next()
2       {
3           ulong i = np_[n_+1];
4
5           if ( i==1 )  return 0;  // current string is last
6
7           if ( dr_[i]==1 )  // direction == "right"
8           {
9               if ( sq_[i] == mx_[i] )   sq_[i] = sq_[i-1] + 1;
10              else                      sq_[i] += 1;
```

```
          positions      direction      delta set
 1:   [ 1 4 7 A ]    [ + + + + ]    1..1..1..1..
 2:   [ 1 4 7 8 ]    [ + + + + ]    1..1..11....
 3:   [ 1 4 7 9 ]    [ + + + - ]    1..1..1.1...
 4:   [ 1 4 5 9 ]    [ + + + - ]    1..11...1...
 5:   [ 1 4 5 8 ]    [ + + + - ]    1..11..1....
 6:   [ 1 4 5 7 ]    [ + + + - ]    1..11.1.....
 7:   [ 1 4 5 6 ]    [ + + + - ]    1..111......
 8:   [ 1 4 5 A ]    [ + + + + ]    1..11....1..
 9:   [ 1 4 6 A ]    [ + + - + ]    1..1.1...1..
10:   [ 1 4 6 7 ]    [ + + - + ]    1..1.11.....
11:   [ 1 4 6 8 ]    [ + + - + ]    1..1.1.1....
12:   [ 1 4 6 9 ]    [ + + - - ]    1..1.1..1...
13:   [ 1 2 6 9 ]    [ + + - - ]    11...1..1...
14:   [ 1 2 6 8 ]    [ + + - - ]    11...1.1....
15:   [ 1 2 6 7 ]    [ + + - - ]    11...11.....
16:   [ 1 2 6 A ]    [ + + - + ]    11...1...1..
17:   [ 1 2 5 A ]    [ + + - + ]    11..1....1..
18:   [ 1 2 5 6 ]    [ + + - + ]    11..11......
19:   [ 1 2 5 7 ]    [ + + - + ]    11..1.1.....
20:   [ 1 2 5 8 ]    [ + + - + ]    11..1..1....
21:   [ 1 2 5 9 ]    [ + + - - ]    11..1...1...
22:   [ 1 2 4 9 ]    [ + + - - ]    11.1....1...
23:   [ 1 2 4 8 ]    [ + + - - ]    11.1...1....
24:   [ 1 2 4 7 ]    [ + + - - ]    11.1..1.....
25:   [ 1 2 4 6 ]    [ + + - - ]    11.1.1......
26:   [ 1 2 4 5 ]    [ + + - - ]    11.11.......
27:   [ 1 2 4 A ]    [ + + - + ]    11.1.....1..
28:   [ 1 2 3 A ]    [ + + - + ]    111......1..
29:   [ 1 2 3 4 ]    [ + + - + ]    1111........
30:   [ 1 2 3 5 ]    [ + + - + ]    111.1.......
31:   [ 1 2 3 6 ]    [ + + - + ]    111..1......
32:   [ 1 2 3 7 ]    [ + + - + ]    111...1.....
33:   [ 1 2 3 8 ]    [ + + - + ]    111....1....
34:   [ 1 2 3 9 ]    [ + + - - ]    111.....1...
35:   [ 1 2 7 9 ]    [ + + + - ]    11.....1.1...
36:   [ 1 2 7 8 ]    [ + + + - ]    11.....11....
37:   [ 1 2 7 A ]    [ + + + + ]    11.....1..1..
38:   [ 1 3 7 A ]    [ + - + + ]    1.1...1..1..
39:   [ 1 3 7 8 ]    [ + - + + ]    1.1...11....
40:   [ 1 3 7 9 ]    [ + - + - ]    1.1...1.1...
41:   [ 1 3 4 9 ]    [ + - + - ]    1.11....1...
42:   [ 1 3 4 8 ]    [ + - + - ]    1.11...1....
43:   [ 1 3 4 7 ]    [ + - + - ]    1.11..1.....
44:   [ 1 3 4 6 ]    [ + - + - ]    1.11.1......
45:   [ 1 3 4 5 ]    [ + - + - ]    1.111.......
46:   [ 1 3 4 A ]    [ + - + + ]    1.11.....1..
47:   [ 1 3 5 A ]    [ + - + + ]    1.1.1....1..
48:   [ 1 3 5 6 ]    [ + - + + ]    1.1.11......
49:   [ 1 3 5 7 ]    [ + - + + ]    1.1.1.1.....
50:   [ 1 3 5 8 ]    [ + - + + ]    1.1.1..1....
51:   [ 1 3 5 9 ]    [ + - + - ]    1.1.1...1...
52:   [ 1 3 6 9 ]    [ + - - - ]    1.1..1..1...
53:   [ 1 3 6 8 ]    [ + - - - ]    1.1..1.1....
54:   [ 1 3 6 7 ]    [ + - - - ]    1.1..11.....
55:   [ 1 3 6 A ]    [ + - - + ]    1.1..1...1..
```

Figure 13.4-B: Gray code for 3-ary trees with 4 internal nodes with all changes being homogeneous. The left column shows the vectors of (one-based) positions, the symbol '`A`' is used for the number 10.

```
11
12              if ( sq_[i] == mx_[i] - 1 )
13              {
14                  np_[i+1] = np_[i];  // can access element n+1
15                  np_[i] = i - 1;
16                  dr_[i] = -1UL;      // "left"
17              }
18          }
19          else
20          {
21              if ( sq_[i] == sq_[i-1] + 1 )
22              {
23                  sq_[i] = mx_[i];
24                  dr_[i] = 1;         // "right"
25                  np_[i+1] = np_[i];  // can access element n+1
26                  np_[i] = i - 1;
27              }
28              else  sq_[i] -= 1;
29          }
30
31          if ( i<n_ )  np_[n_+1] = n_;
32
33          return i - 1;
34      }
35  };
```

The rate of generation is about 97 M/s for 2-ary trees (corresponding to Catalan strings), 120 M/s for 3-ary trees, and 139 M/s with 4-ary trees [FXT: comb/tree-gray-demo.cc].

Chapter 14

Integer partitions

```
 1:    6 ==  6* 1 + 0    + 0    + 0    + 0    + 0     ==   1 + 1 + 1 + 1 + 1 + 1
 2:    6 ==  4* 1 + 1* 2 + 0    + 0    + 0    + 0     ==   1 + 1 + 1 + 1 + 2
 3:    6 ==  2* 1 + 2* 2 + 0    + 0    + 0    + 0     ==   1 + 1 + 2 + 2
 4:    6 ==  0    + 3* 2 + 0    + 0    + 0    + 0     ==   2 + 2 + 2
 5:    6 ==  3* 1 + 0    + 1* 3 + 0    + 0    + 0     ==   1 + 1 + 1 + 3
 6:    6 ==  1* 1 + 1* 2 + 1* 3 + 0    + 0    + 0     ==   1 + 2 + 3
 7:    6 ==  0    + 0    + 2* 3 + 0    + 0    + 0     ==   3 + 3
 8:    6 ==  2* 1 + 0    + 0    + 1* 4 + 0    + 0     ==   1 + 1 + 4
 9:    6 ==  0    + 1* 2 + 0    + 1* 4 + 0    + 0     ==   2 + 4
10:    6 ==  1* 1 + 0    + 0    + 0    + 1* 5 + 0     ==   1 + 5
11:    6 ==  0    + 0    + 0    + 0    + 0    + 1* 6  ==   6
```

Figure 14.0-A: All (eleven) integer partitions of 6.

An integer x is the sum of the positive integers less or equal to itself in various ways. The decompositions into sums of integers are called the *integer partitions* of the number x. Figure 14.0-A shows all integer partitions of $x = 6$.

14.1 Recursive solution of a generalized problem

We can solve a slightly more general problem and find all partitions of a number x with respect to a set $V = \{v_0, v_1, \ldots, v_{n-1}\}$ (where $v_i > 0$), that is all decompositions of the form $x = \sum_{k=0}^{n-1} c_k \cdot v_k$ (where $c_i \geq 0$). The integer partitions are the special case $V = \{1, 2, 3, \ldots, n\}$.

The algorithm to generate the partitions is to assign to the first bucket r_0 an integer multiple of the first set element v_0: $r_0 = c \cdot v_0$ (this has to be done for all $c \geq 0$ for which $r_0 \leq x$). Now set $c_0 = c$. If $r_0 = x$ we already found a partition (consisting of c_0 only), else (if $r_0 < x$) solve the remaining problem where $x' := x - c_0 \cdot v_0$ and $V' := \{v_1, v_2, \ldots, v_{n-1}\}$.

A C++ class for the generation of all partitions is [FXT: `class partition_rec` in comb/partition-rec.h]:

```
1   class partition_rec
2   // Integer partitions of x into supplied values pv[0],...,pv[n-1].
3   // pv[] defaults to [1,2,3,...,x]
4   {
5   public:
6       ulong ct_;  // Number of partitions found so far
7       ulong n_;   // Number of values
8       ulong i_;   // level in iterative search
9
10      long *pv_;  // values into which to partition
11      ulong *pc_; // multipliers for values
12      ulong pci_; // temporary for pc_[i_]
13      long *r_;   // rest
```

```
14      long ri_;   // temporary for r_[i_]
15      long x_;    // value to partition
16
17  public:
18      partition_rec(ulong x, ulong n=0, const ulong *pv=0)
19      {
20          if ( 0==n )  n = x;
21          n_ = n;
22          pv_ = new long[n_+1];
23          if ( pv )  for (ulong j=0; j<n_; ++j)  pv_[j] = pv[j];
24          else       for (ulong j=0; j<n_; ++j)  pv_[j] = j + 1;
25          pc_ = new ulong[n_+1];
26          r_  = new long[n_+1];
27          init(x);
28      }
29
30      void init(ulong x)
31      {
32          x_ = x;
33          ct_ = 0;
34          for (ulong k=0; k<n_; ++k)  pc_[k] = 0;
35          for (ulong k=0; k<n_; ++k)  r_[k] = 0;
36          r_[n_-1] = x_;
37          r_[n_] = x_;
38          i_ = n_ - 1;
39          pci_ = 0;
40          ri_ = x_;
41      }
42
43      ~partition_rec()
44      {
45          delete [] pv_;
46          delete [] pc_;
47          delete [] r_;
48      }
49
50      ulong next();  // generate next partition
51      ulong next_func(ulong i);  // aux
52
53      [--snip--]
54  };
```

The routine to obtains the next partition is given in [FXT: comb/partition-rec.cc], it is actually an iterative version of the algorithm:

```
1   ulong
2   partition_rec::next()
3   {
4       if ( i_>=n_ )  return  n_;
5
6       r_[i_] = ri_;
7       pc_[i_] = pci_;
8       i_ = next_func(i_);
9
10      for (ulong j=0; j<i_; ++j)  pc_[j] = r_[j] = 0;
11
12      ++i_;
13      ri_ = r_[i_] - pv_[i_];
14      pci_ = pc_[i_] + 1;
15
16      return  i_ - 1;  // >=0
17  }
```

```
1   ulong
2   partition_rec::next_func(ulong i)
3   {
4    start:
5       if ( 0!=i )
6       {
7           while ( r_[i]>0 )
8           {
9               pc_[i-1] = 0;
10              r_[i-1] = r_[i];
11              --i;  goto start;   // iteration
12          }
13      }
14      else  // iteration end
```

```
15      {
16          if ( 0!=r_[i] )
17          {
18              long d = r_[i] / pv_[i];
19              r_[i] -= d * pv_[i];
20              pc_[i] = d;
21          }
22      }
23
24      if ( 0==r_[i] )  // valid partition found
25      {
26          ++ct_;
27          return i;
28      }
29
30      ++i;
31      if ( i>=n_ )  return n_;  // search finished
32
33      r_[i] -= pv_[i];
34      ++pc_[i];
35
36      goto start;  // iteration
37  }
```

The routines can easily adapted to the generation of partitions satisfying certain restrictions, for example, partitions into unequal parts (that is, $c_i \leq 1$).

The listing shown in figure 14.0-A can be generated with [FXT: comb/partition-rec-demo.cc]:

```
1   void
2   print_part(ulong n)
3   {
4           partition pp(n);
5           ulong ct = 0;
6           while ( pp.next() < n )
7           {
8               cout << "    #" << setw(2) << ct << ": ";
9               cout << setw(4) << pp.x_ << " == ";
10              pp.print2();
11              cout << "  ==  ";
12              pp.print();
13              cout << endl;
14              ++ct;
15          }
16          cout << " " << n << ":  ct=" << ct << endl;
17          cout << endl;
18  }
```

The 190,569,292 partitions of 100 are generated in less than 11 seconds, corresponding to a rate of about 18 million partitions per second.

14.2 Iterative algorithm

An iterative implementation for the special case $V = \{1, 2, 3, \ldots, n\}$ (the integer partitions) is given in [FXT: `class partition` in comb/partition.h]:

```
1   class partition
2   {
3   public:
4       ulong *c_;  // partition:  c[1]* 1 + c[2]* 2 + ... + c[n]* n == n
5       ulong *s_;  // cumulative sums:  s[j+1] = c[1]* 1 + c[2]* 2 + ... + c[j]* j
6       ulong n_;   // partitions of n
7
8   public:
9       partition(ulong n)
10      {
11          n_ = n;
12          c_ = new ulong[n+1];
13          s_ = new ulong[n+1];
14          s_[0] = 0;  // unused
15          c_[0] = 0;  // unused
```

```
16          first();
17      }
18
19      ~partition()
20      {
21          delete [] c_;
22          delete [] s_;
23      }
24
25      void first()
26      {
27          c_[1] = n_;
28          for (ulong i=2; i<=n_; i++)  { c_[i] = 0; }
29          s_[1] = 0;
30          for (ulong i=2; i<=n_; i++)  { s_[i] = n_; }
31      }
32
33      void last()
34      {
35          for (ulong i=1; i<n_; i++)  { c_[i] = 0; }
36          c_[n_] = 1;
37          for (ulong i=1; i<n_; i++)  { s_[i] = 0; }
38          // s_[n_+1] = n_;  // unused (and out of bounds)
39      }
```

To obtain the next partition, find the smallest index $i \geq 2$ so that $[c_1, c_2, \ldots, c_{i-1}, c_i]$ can be replaced by $[z, 0, 0, \ldots, 0, c_i + 1]$ where $z \geq 0$. The index i (and z) is determined using cumulative sums. The partitions are generated in the same order as shown in figure 14.0-A. The algorithm was given (2006) by Torsten Finke [priv.comm.].

```
1       bool next()
2       {
3           if ( c_[n_]!=0 )  return false;  // last == 1* n (c[n]==1)
4
5           // Find first coefficient c[i], i>=2 that can be increased:
6           ulong i = 2;
7           while ( s_[i]<i )  ++i;
8
9           ++c_[i];
10          s_[i] -= i;
11          ulong z = s_[i];
12          // Now set c[1], c[2], ..., c[i-1] to the first partition
13          // of z into i-1 parts, i.e. set to  z, 0, 0, ..., 0:
14          while ( --i > 1 )
15          {
16              s_[i] = z;
17              c_[i] = 0;
18          }
19          c_[1] = z;  // z* 1 == z
20          // s_[1] unused
21
22          return true;
23      }
```

The preceding partition can be computed as follows:

```
1       bool prev()
2       {
3           if ( c_[1]==n_ )  return false;  // first == n* 1 (c[1]==n)
4
5           // Find first nonzero coefficient c[i] where i>=2:
6           ulong i = 2;
7           while ( c_[i]==0 )  ++i;
8
9           --c_[i];
10          s_[i] += i;
11          ulong z = s_[i];
12          // Now set c[1], c[2], ..., c[i-1] to the last partition
13          // of z into i-1 parts:
14          while ( --i > 1  )
15          {
16              ulong q = (z>=i ? z/i : 0); // == z/i;
17              c_[i] = q;
18              s_[i+1] = z;
19              z -= q*i;
20          }
21          c_[1] = z;
```

```
22          s_[2] = z;
23          // s_[1] unused
24
25          return true;
26      }
27
28    [--snip--]
29  };
```

Note that divisions which result in $q = 0$ are avoided, leading to a small speedup. The program [FXT: comb/partition-demo.cc] demonstrates the usage of the class. More than 140 million partitions per second are generated, about 66 million when going backward.

14.3 Partitions into m parts

```
 1:  1 1 1 1 1 1 1 1 1 1 9          12:  1 1 1 1 1 1 1 2 2 3 5
 2:  1 1 1 1 1 1 1 1 1 2 8          13:  1 1 1 1 1 1 1 2 2 4 4
 3:  1 1 1 1 1 1 1 1 1 3 7          14:  1 1 1 1 1 1 1 2 3 3 4
 4:  1 1 1 1 1 1 1 1 1 4 6          15:  1 1 1 1 1 1 1 3 3 3 3
 5:  1 1 1 1 1 1 1 1 1 5 5          16:  1 1 1 1 1 1 2 2 2 2 5
 6:  1 1 1 1 1 1 1 1 2 2 7          17:  1 1 1 1 1 1 2 2 2 3 4
 7:  1 1 1 1 1 1 1 1 2 3 6          18:  1 1 1 1 1 1 2 2 3 3 3
 8:  1 1 1 1 1 1 1 1 2 4 5          19:  1 1 1 1 1 2 2 2 2 2 4
 9:  1 1 1 1 1 1 1 1 3 3 5          20:  1 1 1 1 1 2 2 2 2 3 3
10:  1 1 1 1 1 1 1 1 3 4 4          21:  1 1 1 1 2 2 2 2 2 2 3
11:  1 1 1 1 1 1 1 2 2 2 6          22:  1 1 1 2 2 2 2 2 2 2 2
```

Figure 14.3-A: The 22 partitions of 19 into 11 parts in lexicographic order.

An algorithm for the generation of all partitions of n into m parts is given in [104, p.106] (method ascribed to Hindenburg):

> The initial partition contains $m-1$ units and the element $n-m+1$. To obtain a new partition from a given one, pass over the elements of the latter from right to left, stopping at the first element f which is less, by at least two units, than the final element [...]. Without altering any element at the left of f, write $f+1$ in place of f and every element to the right of f with the exception of the final element, in whose place is written the number which when added to all the other new elements gives the sum n. The process to obtain partitions stops when we reach one in which no part is less than the final part by at least two units.

Figure 14.3-A shows the partitions of 19 into 11 parts. The data was generated with the program [FXT: comb/mpartition-demo.cc].

An efficient implementation is given as [FXT: `class mpartition` in comb/mpartition.h]:

```
1   class mpartition
2   // Integer partitions of n into m parts
3   {
4   public:
5       ulong *x_;  // partition: x[1]+x[2]+...+x[m] = n
6       ulong *s_;  // aux: cumulative sums of x[]  (s[0]=0)
7       ulong n_;   // integer partitions of n  (must have n>0)
8       ulong m_;   // ... into m parts  (must have 0<m<=n)
9
10  public:
11      mpartition(ulong n, ulong m)
12          : n_(n), m_(m)
13      {
14          x_ = new ulong [m_+1];
15          s_ = new ulong [m_+1];
16          init();
17      }
18
19      ~mpartition()
20      {
21          delete [] x_;
22          delete [] s_;
```

```
23      }
24
25      const ulong *data() const  { return x_+1; }
26
27      void init()
28      {
29          x_[0] = 0;
30          for (ulong k=1; k<m_; ++k)  x_[k] = 1;
31          x_[m_] = n_ - m_ + 1;
32          ulong s = 0;
33          for (ulong k=0; k<=m_; ++k)  { s+=x_[k]; s_[k]=s; }
34      }
35
36      bool next()
37      {
38          ulong u = x_[m_];  // last element
39          ulong k = m_;
40          while ( --k )  { if ( x_[k]+2<=u ) break; }
41
42          if ( k==0 )  return  false;
43
44          ulong f = x_[k] + 1;
45          ulong s = s_[k-1];
46          while ( k < m_ )
47          {
48              x_[k] = f;
49              s += f;
50              s_[k] = s;
51              ++k;
52          }
53          x_[m_] = n_ - s_[m_-1];
54          // s_[m_] = n_;  // unchanged
55
56          return  true;
57      }
58  };
```

The auxiliary array of cumulative sums allows the recalculation of the final element without rescanning more than the elements just changed. About 105 million partitions per second can be generated.

A (complicated) construction for a Gray code for integer partitions is given in [220].

14.4 The number of integer partitions

n:	P_n	n:	P_n	n:	P_n	n:	P_n	n:	P_n
1:	1	11:	56	21:	792	31:	6842	41:	44583
2:	2	12:	77	22:	1002	32:	8349	42:	53174
3:	3	13:	101	23:	1255	33:	10143	43:	63261
4:	5	14:	135	24:	1575	34:	12310	44:	75175
5:	7	15:	176	25:	1958	35:	14883	45:	89134
6:	11	16:	231	26:	2436	36:	17977	46:	105558
7:	15	17:	297	27:	3010	37:	21637	47:	124754
8:	22	18:	385	28:	3718	38:	26015	48:	147273
9:	30	19:	490	29:	4565	39:	31185	49:	173525
10:	42	20:	627	30:	5604	40:	37338	50:	204226

Figure 14.4-A: The number of integer partitions of n for $n \leq 50$.

The total number of integer partitions of n is sequence A000041 of [240], the values for $1 \leq x \leq 50$ are shown in figure 14.4-A. If we denote the number of partitions of n into exactly m parts by $P(n,m)$ then

$$P(n,m) \quad = \quad P(n-1,m-1) + P(n-m,m) \tag{14.4-1}$$

n:	P(n)	P(n,m) for m = 1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1:	1	1															
2:	2	1	1														
3:	3	1	1	1													
4:	5	1	2	1	1												
5:	7	1	2	2	1	1											
6:	11	1	3	3	2	1	1										
7:	15	1	3	4	3	2	1	1									
8:	22	1	4	5	5	3	2	1	1								
9:	30	1	4	7	6	5	3	2	1	1							
10:	42	1	5	8	9	7	5	3	2	1	1						
11:	56	1	5	10	11	10	7	5	3	2	1	1					
12:	77	1	6	12	15	13	11	7	5	3	2	1	1				
13:	101	1	6	14	18	18	14	11	7	5	3	2	1	1			
14:	135	1	7	16	23	23	20	15	11	7	5	3	2	1	1		
15:	176	1	7	19	27	30	26	21	15	11	7	5	3	2	1	1	
16:	231	1	8	21	34	37	35	28	22	15	11	7	5	3	2	1	1

Figure 14.4-B: Numbers $P(n,m)$ of partitions of n into m parts.

were we set $P(0,0) = 1$. We obviously have $P_n = \sum_{m=1}^{n} P(n,m)$. Figure 14.4-B shows $P(n,m)$ for $n \le 16$, it was created with the program [FXT: comb/num-partitions-demo.cc]. We note that the number of partitions into m parts equals the number of partitions with maximal part equal to m. This can easily be seen by drawing a diagram and its transposed as follows (for the partition $5+2+2+1$ of 10):

```
   43111      5221
5  XXXXX    4 XXXX
2  XX       3 XXX
2  XX       1 X
1  X        1 X
            1 X
```

Thereby any partition with maximal part m (here 5) corresponds to a partition into exactly m parts. The generating function for the partitions into exactly m parts is

$$\sum_{n=1}^{\infty} P(n,m)\, x^n \;=\; \frac{x^m}{\prod_{k=1}^{m}(1-x^k)} \tag{14.4-2}$$

For example, the row for $m = 3$ in figure 14.4-B corresponds to the power series

```
? m=3; (x^m/prod(k=1,m,1-x^k)+O(x^17))
  x^3 + x^4 + 2*x^5 + 3*x^6 + 4*x^7 + 5*x^8 + 7*x^9 + 8*x^10 + \
  10*x^11 + 12*x^12 + 14*x^13 + 16*x^14 + 19*x^15 + 21*x^16 + O(x^17)
```

The generating function for the number P_n of integer partitions of n can be given as [115, p.357]

$$\sum_{n=0}^{\infty} P_n\, x^n \;=\; \frac{1}{\prod_{n=1}^{\infty}(1-x^n)} \;=:\; \frac{1}{\eta(x)} \tag{14.4-3a}$$

Then we have the *pentagonal number theorem*, due to Euler, see [34]:

$$\eta(x) \;=\; \sum_{n=-\infty}^{+\infty} (-1)^n \left(x^{n(3n-1)/2}\right) \;=\; 1 + \sum_{n=1}^{\infty} (-1)^n \left(x^{n(3n-1)/2} + x^{n(3n+1)/2}\right) \tag{14.4-3b}$$

and

$$\frac{1}{\eta(x)} \;=\; 1 + \sum_{n=1}^{\infty} \frac{x^n}{\prod_{k=1}^{n}(1-x^k)} \tag{14.4-3c}$$

$$\;=\; 1 + \sum_{n=1}^{\infty} \frac{x^{n^2}}{\left[\prod_{k=1}^{n}(1-x^k)\right]^2} \tag{14.4-3d}$$

```
? N=10;x=t+O(t^N);
? 1/prod(k=1,N,1-x^k)
  1 + t + 2*t^2 + 3*t^3 + 5*t^4 + 7*t^5 + 11*t^6 + 15*t^7 + 22*t^8 + 30*t^9 + O(t^10)
\\ == 1+sum(n=1,N,x^n/prod(j=1,n,1-x^j))
\\ == 1+sum(n=1,N,x^(n^2)/(prod(k=1,n,(1-x^k))^2))
\\ == 1/(1+sum(n=1,N,(-1)^n*(x^(n*(3*n-1)/2)+x^(n*(3*n+1)/2))))
```

Relation 14.4-3c is the special case $a_n = x^n$ (and $N \to \infty$) of [199, p.83]

$$\frac{1}{\prod_{n=1}^{N}(1-a_n)} \quad = \quad 1 + \sum_{n=1}^{N} \frac{a_n}{\prod_{k=1}^{n}(1-a_k)} \tag{14.4-4}$$

For the number D_n of partitions of n into distinct parts we have the generating function

$$\eta_+(x) \quad := \quad \prod_{n=1}^{\infty}(1+x^n) \;=\; \sum_{n=0}^{\infty} D_n\, x^n \tag{14.4-5}$$

We have (the number of partitions into distinct parts equals the number of partitions into odd parts)

$$\eta_+(x) \quad = \quad \frac{\eta(x^2)}{\eta(x)} \;=\; \frac{1}{\prod_{k=1}^{\infty}(1-x^{2\,k-1})} \tag{14.4-6}$$

Thereby relation 14.4-3b allows fast computation of $\eta_+(x)$. The sequence of coefficients D_n is entry A000009 of [240]:

```
[1, 1, 1, 2, 2, 3, 4, 5, 6, 8, 10, 12, 15, 18, 22, 27, 32, 38, 46, 54,
 64, 76, 89, 104, 122, 142, 165, 192, 222, 256, ...]
```

The number of partitions where at most r elements of each partition are equal has the generating function

$$\prod_{n=1}^{\infty}\left(1+x+x^2+\ldots+x^r\right) \quad = \quad \frac{\eta(x^{r+1})}{\eta(x)} \;=\; \frac{1}{\prod_{k\neq 0 \bmod r+1}(1-x^k)} \tag{14.4-7a}$$

The second equality tells us that the number of such partitions equals the number of partitions into parts not divisible by $r+1$.

Chapter 15

Set partitions

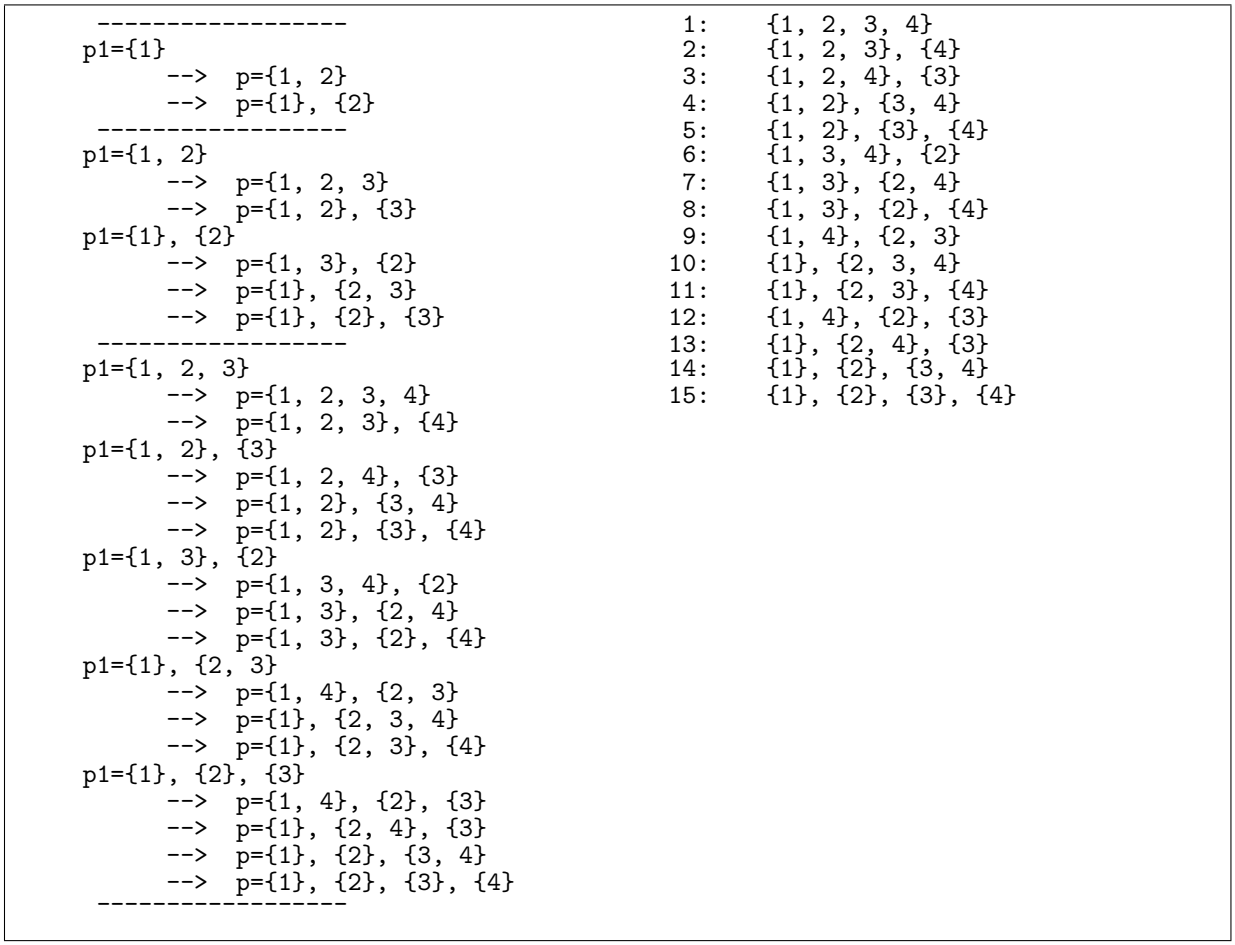

Figure 15.0-A: Recursive construction of the set partitions of the 4-element set $S_4 = \{1, 2, 3, 4\}$ (left). The resulting list of all set partitions of is shown on the right.

For a set of n elements, say $S_n := \{1, 2, \ldots, n\}$, a *set partition* is a set $P = \{s_1, s_2, \ldots, s_k\}$ of non-empty subsets s_i of S_n whose intersection is empty and whose union equals S.

For example, there are 5 set partitions of the set $S_3 = \{1, 2, 3\}$:

```
1:  { {1, 2, 3} }
2:  { {1, 2}, {3} }
3:  { {1, 3}, {2} }
4:  { {1}, {2, 3} }
5:  { {1}, {2}, {3} }
```

The following sets are not set partitions of S_3:

```
{ {1, 2, 3}, {1} }   // intersection not empty
{ {1}, {3} }         // union does not contain 2
```

As the order of elements in a set does not matter we sort them in ascending order. For a set of sets we order the sets in ascending order of the first elements.

Recursive generation

We write Z_n for the list of all set partitions of the n-element set S_n. In order to generate Z_n we observe that with a complete list Z_{n-1} of partitions of the set S_{n-1} we can generate the elements of Z_n in the following way: For each element (set partition) $P \in Z_{n-1}$, create set partitions of S_n by appending the element n to the first, second, ..., last subset and one more by appending the set n as the last subset.

For example, the partition $\{\{1, 2\}, \{3, 4\}\} \in Z_4$ leads to 3 partitions of S_5:

```
P = { {1, 2}, {3, 4} }
      -->  { {1, 2, 5}, {3, 4} }
      -->  { {1, 2}, {3, 4, 5} }
      -->  { {1, 2}, {3, 4}, {5} }
```

Now we start with the only partition of the one-element set, $\{\{1\}\}$, and apply the described step $n - 1$ times. The construction is shown in the left column of figure 15.0-A, the right column shows all set partitions for $n = 5$.

15.1 The number of set partitions: Stirling set numbers and Bell numbers

```
 n   b=   B(n)    k: 1     2     3     4     5     6     7     8     9    10
 1:  b=      1       1
 2:  b=      2       1     1
 3:  b=      5       1     3     1
 4:  b=     15       1     7     6     1
 5:  b=     52       1    15    25    10     1
 6:  b=    203       1    31    90    65    15     1
 7:  b=    877       1    63   301   350   140    21     1
 8:  b=   4140       1   127   966  1701  1050   266    28     1
 9:  b=  21147       1   255  3025  7770  6951  2646   462    36     1
10:  b= 115975       1   511  9330 34105 42525 22827  5880   750    45     1
```

Figure 15.1-A: Stirling numbers of the second kind and Bell numbers.

The sequence of numbers of set partitions of S_n for $n \geq 1$ is 1, 2, 5, 15, 52, 203, 877, ... These are the *Bell numbers*, sequence A000110 of [240]. The Bell numbers can be computed using the observation that in every step of the computation a partition of S_{x-1} into k subsets leads to $k + 1$ partitions of S_x. Of these partitions k contain k subsets and one contains $k + 1$ subsets.

If we write the number of partitions of S_n into k subsets as a triangular array, then the entry can be computed as the sum of its upper left neighbor plus k times its upper neighbor. We start with a single one in the row for $n = 1$. The scheme is shown in figure 15.1-A. The numbers are the *Stirling numbers of the second kind* (or *Stirling set numbers*), see sequence A008277 of [240].

The sum over all elements of row n gives the n-th Bell number. Another way of computing the Bell numbers is given in section 3.8.2 on page 139. The array shown can be generated with the program [FXT: comb/stirling2-demo.cc].

We further have the recursion

$$B_{n+1} \quad = \quad \sum_{k=0}^{n} \binom{n}{k} B_k \tag{15.1-1}$$

As pari/gp code:

```
? N=11; v=vector(N); v[1]=1;
? for (j=2, N, v[j]=sum(k=1, j-1, binomial(j-2,k-1)*v[k])); v
  [1, 1, 2, 5, 15, 52, 203, 877, 4140, 21147, 115975]
```

15.1.1 Generating functions

The ordinary generating function for the Bell numbers can be given as

$$\sum_{n=0}^{\infty} B_n\, x^n \quad = \quad \sum_{k=0}^{\infty} \frac{x^k}{\prod_{j=1}^{k} (1 - j\, x)} \quad = \quad 1 + x + 2\, x^2 + 5\, x^3 + 15\, x^4 + 52\, x^5 + \ldots \tag{15.1-2}$$

The exponential generating function (EGF) is

$$\exp(\exp(x) - 1) \quad = \quad \sum_{n=0}^{\infty} B_n\, \frac{x^n}{n!} \tag{15.1-3}$$

```
? sum(k=0,11,x^k/prod(j=1,k,1-j*x))+O(x^8)  \\ OGF
  1 + x + 2*x^2 + 5*x^3 + 15*x^4 + 52*x^5 + 203*x^6 + 877*x^7 + O(x^8)
? serlaplace(exp(exp(x)-1))                 \\ EGF
  1 + x + 2*x^2 + 5*x^3 + 15*x^4 + 52*x^5 + 203*x^6 + 877*x^7 + 4140*x^8 + ...
```

Dobinski's formula for the Bell numbers is

$$B_n \quad = \quad \frac{1}{e} \sum_{k=1}^{\infty} \frac{n^k}{k!} \tag{15.1-4}$$

The array of Stirling numbers shown in figure 15.1-A can also be computed in polynomial form by setting $B_0(x) = 1$ and

$$B_{n+1}(x) \quad = \quad x\, (B_n'(x) + B_x(x)) \tag{15.1-5}$$

The coefficients of $B_n(x)$ are the Stirling numbers and $B_n(1) = B_n$:

```
? B=1; for(k=1,6, B=x*(deriv(B)+B); print(subst(B,x,1),": ",B))
   1: x
   2: x^2 + x
   5: x^3 + 3*x^2 + x
  15: x^4 + 6*x^3 + 7*x^2 + x
  52: x^5 + 10*x^4 + 25*x^3 + 15*x^2 + x
 203: x^6 + 15*x^5 + 65*x^4 + 90*x^3 + 31*x^2 + x
```

The polynomials are called *Bell polynomials*, see [268].

15.1.2 Set partitions of a given type

We say a set partition of the n-element set is of *type* $C = [c_1, c_2, c_3, \ldots, c_n]$ if it has c_1 one-element sets, c_2 two-element sets, c_3 three-element sets, and so on. Define

$$L(z) \quad = \quad \sum_{k=1}^{\infty} \frac{c_k\, z^k}{k!} \tag{15.1-6}$$

then

$$\exp(L(z)) \quad = \quad \sum_{n=0}^{\infty} \left[\sum_{C} \left(Z_{n,C} \prod t_k^{c_k} \right) \right] \frac{z^n}{n!} \tag{15.1-7}$$

where $Z_{n,C}$ is the number of set partitions of the n-element set with type C.

```
? n=8;R=O(z^(n+1));
? L=sum(k=1,8,eval(Str("t"k))*z^k/k!)+R
  t1*z + 1/2*t2*z^2 + 1/6*t3*z^3 + 1/24*t4*z^4 +  [...]  + 1/40320*t8*z^8 + O(z^9)
? serlaplace(exp(L))
 1
 + t1 *z
 + (t1^2 + t2) *z^2
 + (t1^3 + 3*t2*t1 + t3) *z^3
 + (t1^4 + 6*t2*t1^2 + 4*t3*t1 + 3*t2^2 + t4) *z^4
 + (t1^5 + 10*t2*t1^3 + 10*t3*t1^2 + 15*t1*t2^2 + 5*t1*t4 + 10*t3*t2 + t5) *z^5
 + (t1^6 + 15*t2*t1^4 + 20*t3*t1^3 +  [...]  + 15*t2^3 + 15*t4*t2 + 10*t3^2 + t6) *z^6
 + (t1^7 + 21*t2*t1^5 + 35*t3*t1^4 +  [...]  + 105*t3*t2^2 + 21*t5*t2 + 35*t4*t3 + t7) *z^7
 + (t1^8 + 28*t2*t1^6 + 56*t3*t1^5 +  [...]  + 28*t6*t2 + 56*t5*t3 + 35*t4^2 + t8) *z^8
 + O(z^9)
```

Specializations give generating functions for set partitions with certain restrictions. For example, the EGF for the partitions without sets of size one is (set $t_1 = 0$ and $t_k = 1$ for $k \neq 1$) $\exp(\exp(z) - 1 - z)$, see entry A000296 of [240]. Section 10.15.1.2 on page 283 gives a similar construction for the EGF for permutations of prescribed cycle type.

15.2 Generation in minimal-change order

```
------------------
P={1}
  --> {1, 2}
  --> {1}, {2}

------------------
P={1, 2}
  --> {1, 2, 3}
  --> {1, 2}, {3}

P={1}, {2}
  -->{1}, {2}, {3}
  -->{1}, {2, 3}
  -->{1, 3}, {2}
```

```
------------------
P={1, 2, 3}
  --> {1, 2, 3, 4}
  --> {1, 2, 3}, {4}

P={1, 2}, {3}
  --> {1, 2}, {3}, {4}
  --> {1, 2}, {3, 4}
  --> {1, 2, 4}, {3}

P={1}, {2}, {3}
  --> {1, 4}, {2}, {3}
  --> {1}, {2, 4}, {3}
  --> {1}, {2}, {3, 4}
  --> {1}, {2}, {3}, {4}

P={1}, {2, 3}
  --> {1}, {2, 3}, {4}
  --> {1}, {2, 3, 4}
  --> {1, 4}, {2, 3}

P={1, 3}, {2}
  --> {1, 3, 4}, {2}
  --> {1, 3}, {2, 4}
  --> {1, 3}, {2}, {4}
```

```
setpart(4)==
{1, 2, 3, 4}
{1, 2, 3}, {4}
{1, 2}, {3}, {4}
{1, 2}, {3, 4}
{1, 2, 4}, {3}
{1, 4}, {2}, {3}
{1}, {2, 4}, {3}
{1}, {2}, {3, 4}
{1}, {2}, {3}, {4}
{1}, {2, 3}, {4}
{1}, {2, 3, 4}
{1, 4}, {2, 3}
{1, 3, 4}, {2}
{1, 3}, {2, 4}
{1, 3}, {2}, {4}
```

Figure 15.2-A: Construction of a Gray code for set partitions as an interleaving process.

A modified version of the recursive construction generates the set partitions in a minimal-change order. We can generate the 'incremented' partitions in two orders, forward (left to right)

```
P = { {1, 2}, {3, 4} }
      -->  { {1, 2, 5}, {3, 4} }
      -->  { {1, 2}, {3, 4, 5} }
      -->  { {1, 2}, {3, 4}, {5} }
```

or backward (right to left)

```
 1:  {1, 2, 3, 4}                  1:  {1}, {2}, {3}, {4}
 2:  {1, 2, 3}, {4}                2:  {1}, {2}, {3, 4}
 3:  {1, 2}, {3}, {4}              3:  {1}, {2, 4}, {3}
 4:  {1, 2}, {3, 4}                4:  {1, 4}, {2}, {3}
 5:  {1, 2, 4}, {3}                5:  {1, 4}, {2, 3}
 6:  {1, 4}, {2}, {3}              6:  {1}, {2, 3, 4}
 7:  {1}, {2, 4}, {3}              7:  {1}, {2, 3}, {4}
 8:  {1}, {2}, {3, 4}              8:  {1, 3}, {2}, {4}
 9:  {1}, {2}, {3}, {4}            9:  {1, 3}, {2, 4}
10:  {1}, {2, 3}, {4}             10:  {1, 3, 4}, {2}
11:  {1}, {2, 3, 4}               11:  {1, 2, 3, 4}
12:  {1, 4}, {2, 3}               12:  {1, 2, 3}, {4}
13:  {1, 3, 4}, {2}               13:  {1, 2}, {3}, {4}
14:  {1, 3}, {2, 4}               14:  {1, 2}, {3, 4}
15:  {1, 3}, {2}, {4}             15:  {1, 2, 4}, {3}
```

Figure 15.2-B: Set partitions of $S_4 = \{1, 2, 3, 4\}$ in two different minimal-change orders.

```
P = { {1, 2}, {3, 4} }
      -->  { {1, 2}, {3, 4}, {5} }
      -->  { {1, 2}, {3, 4, 5} }
      -->  { {1, 2, 5}, {3, 4} }
```

The resulting process of interleaving elements is shown in figure 15.2-A. The method is similar to Trotter's construction for permutations, see figure 10.7-B on page 254. If we change the direction with every subset that is to be incremented, we obtain the minimal-change order shown in figure 15.2-B for $n = 4$. The left column is obtained by starting with the forward direction in each step of the recursion, the right by starting with the backward direction. The lists can be computed with [FXT: comb/setpartition-demo.cc].

The C++ class [FXT: `class set_partition` in comb/setpartition.h] stores the list in an array of signed characters. The stored value is negated if the element is the last in the subset. The work involved with the creation of Z_n is proportional to $\sum_{k=1}^{n} k\, B_k$ where B_k is the k-th Bell number.

The parameter `xdr` of the constructor determines the order in which the partitions are being created:

```
1   class set_partition
2   // Set partitions of the set {1,2,3,...,n}
3   // By default in minimal-change order
4   {
5   public:
6       ulong n_;   // Number of elements of set (set = {1,2,3,...,n})
7       int *p_;    // p[] contains set partitions of length 1,2,3,...,n
8       int **pp_;  // pp[k] points to start of set partition k
9       int *ns_;   // ns[k] Number of Sets in set partition k
10      int *as_;   // element k attached At Set (0<=as[k]<=k) of set(k-1)
11      int *d_;    // direction with recursion (+1 or -1)
12      int *x_;    // current set partition (==pp[n])
13      bool xdr_;  // whether to change direction in recursion (==> minimal-change order)
14      int dr0_;   // dr0: starting direction in each recursive step:
15      //   dr0=+1  ==> start with partition  {{1,2,3,...,n}}
16      //   dr0=-1  ==> start with partition  {{1},{2},{3},...,{n}}}
17
18  public:
19      set_partition(ulong n, bool xdr=true, int dr0=+1)
20          : n_(n)
21      {
22          ulong np = (n_*(n_+1))/2;  // == \sum_{k=1}^{n}{k}
23          p_ = new int[np];
24
25          pp_ = new int *[n_+1];
26          pp_[0] = 0;  // unused
27          pp_[1] = p_;
28          for (ulong k=2; k<=n_; ++k)  pp_[k] = pp_[k-1] + (k-1);
29
30          ns_ = new int[n_+1];
31          as_ = new int[n_+1];
32          d_ = new int[n_+1];
33          x_ = pp_[n_];
```

```
34
35          init(xdr, dr0);
36      }
37     [--snip--]  // destructor
38
39      bool next()  { return next_rec(n_); }
40
41      const int* data()  const  { return x_; }
42
43      ulong print()  const
44      // Print current set partition
45      // Return number of chars printed
46      { return print_p(n_); }
47
48      ulong print_p(ulong k)  const;
49      void print_internal()  const;  // print internal state
50
51  protected:
52     [--snip--]  // internal methods
53  };
```

The actual work is done by the methods `next_rec()` and `cp_append()` [FXT: comb/setpartition.cc]:

```
1   int
2   set_partition::cp_append(const int *src, int *dst, ulong k, ulong a)
3   // Copy partition in src[0,...,k-2] to dst[0,...,k-1]
4   // append element k at subset a (a>=0)
5   // Return number of sets in created partition.
6   {
7       ulong ct = 0;
8       for (ulong j=0; j<k-1; ++j)
9       {
10          int e = src[j];
11          if ( e > 0 )  dst[j] = e;
12          else
13          {
14              if ( a==ct )  { dst[j]=-e; ++dst; dst[j]=-k; }
15              else  dst[j] = e;
16              ++ct;
17          }
18      }
19      if ( a>=ct )  { dst[k-1] = -k; ++ct; }
20
21      return ct;
22  }
```

```
1   int
2   set_partition::next_rec(ulong k)
3   // Update partition in level k from partition in level k-1  (k<=n)
4   // Return number of sets in created partition
5   {
6       if ( k<=1 )  return 0;  // current is last
7
8       int d = d_[k];
9       int as = as_[k] + d;
10      bool ovq = ( (d>0) ? (as>ns_[k-1]) : (as<0) );
11      if ( ovq )  // have to recurse
12      {
13          ulong ns1 = next_rec(k-1);
14          if ( 0==ns1 )  return 0;
15
16          d = ( xdr_ ? -d : dr0_ );
17          d_[k] = d;
18
19          as = ( (d>0) ? 0 : ns_[k-1] );
20      }
21      as_[k] = as;
22
23      ulong ns = cp_append(pp_[k-1], pp_[k], k, as);
24      ns_[k] = ns;
25      return ns;
26  }
```

The partitions are stored as an array of signed integers that are greater than zero and smaller or equal to n. A negative value indicates that it is the last of the subset. For example, the set partitions of S_4 together with their 'signed value' representations are shown in figure 15.2-C. The array `as[]` contains

```
 1:  as[ 0 0 0 0 ]    x[ +1 +2 +3 -4 ]    {1, 2, 3, 4}
 2:  as[ 0 0 0 1 ]    x[ +1 +2 -3 -4 ]    {1, 2, 3}, {4}
 3:  as[ 0 0 1 0 ]    x[ +1 +2 -4 -3 ]    {1, 2, 4}, {3}
 4:  as[ 0 0 1 1 ]    x[ +1 -2 +3 -4 ]    {1, 2}, {3, 4}
 5:  as[ 0 0 1 2 ]    x[ +1 -2 -3 -4 ]    {1, 2}, {3}, {4}
 6:  as[ 0 1 0 0 ]    x[ +1 +3 -4 -2 ]    {1, 3, 4}, {2}
 7:  as[ 0 1 0 1 ]    x[ +1 -3 +2 -4 ]    {1, 3}, {2, 4}
 8:  as[ 0 1 0 2 ]    x[ +1 -3 -2 -4 ]    {1, 3}, {2}, {4}
 9:  as[ 0 1 1 0 ]    x[ +1 -4 +2 -3 ]    {1, 4}, {2, 3}
10:  as[ 0 1 1 1 ]    x[ -1 +2 +3 -4 ]    {1}, {2, 3, 4}
11:  as[ 0 1 1 2 ]    x[ -1 +2 -3 -4 ]    {1}, {2, 3}, {4}
12:  as[ 0 1 2 0 ]    x[ +1 -4 -2 -3 ]    {1, 4}, {2}, {3}
13:  as[ 0 1 2 1 ]    x[ -1 +2 -4 -3 ]    {1}, {2, 4}, {3}
14:  as[ 0 1 2 2 ]    x[ -1 -2 +3 -4 ]    {1}, {2}, {3, 4}
15:  as[ 0 1 2 3 ]    x[ -1 -2 -3 -4 ]    {1}, {2}, {3}, {4}
```

Figure 15.2-C: The partitions of the set $S_4 = \{1, 2, 3, 4\}$ together with the internal representations: the 'signed value' array `x[]` and the 'attachment' array `as[]`.

a restricted growth string (RGS) with the condition $a_j \leq 1 + \max_{i<j}(a_i)$. A different sort of RGS is described in section 13.2 on page 315.

The copying is the performance bottleneck of the algorithm. Therefore 'only' about 11 million partitions are generated per second. An $O(1)$ algorithm for the Gray code starting with all elements in one set is given in [161].

15.3 Restricted growth strings

For some applications the restricted growth strings (RGS) may suffice. We give an implementation for their generation and algorithms to generate two classes of generalized RGS that contain the RGS for set partitions as a special case.

15.3.1 RGS for set partitions in minimal-change order

The C++ implementation [FXT: `class set_partition_rgs` in comb/setpartition-rgs.h] generates the RGS for set partitions in lexicographic or minimal-change order. The RGS are updated by the recursive routine [FXT: comb/setpartition-rgs.cc]:

```
1   int
2   set_partition_rgs::next_rec(ulong k)
3   {
4       if ( k<=1 )  return 0;  // current is last
5
6       int d = d_[k];
7       int as = as_[k] + d;
8       bool ovq = ( (d>0) ? (as>ns_[k-1]) : (as<0) );
9       if ( ovq )  // have to recurse
10      {
11          if ( 0==next_rec(k-1) )  return 0;
12
13          d = ( xdr_ ? -d : dr0_ );
14          d_[k] = d;
15
16          as = ( (d>0) ? 0 : ns_[k-1] );
17      }
18      as_[k] = as;
19
20      ulong ns = ns_[k] = max2(ns_[k-1], as_[k]+1);
21      return ns;
22  }
```

An iterative version of the update routine is

```
1       bool next()
2       {
3           ulong k = n_;
4           while ( (ulong)(as_[k] + d_[k]) > (ulong)ns_[k-1] ) // <0 or >max
5           {
6               if ( --k <= 1 )  return 0;
7           }
8
9           as_[k] += d_[k];
10          ns_[k] = max2(ns_[k-1], as_[k]+1);
11
12          while ( ++k<=n_ )
13          {
14              ulong d = d_[k];
15              d = ( xdr_ ? -d : dr0_ );
16              d_[k] = d;
17              ulong as = ( ((int)d>0) ? 0 : ns_[k-1] );
18              as_[k] = as;
19              ns_[k] = max2(ns_[k-1], as_[k]+1);
20          }
21
22          return 1;
23      }
```

It is activated (by default) via the define

```
#ifdef SETPART_RGS_ITERATIVE
```

A program that shows the usage of the class is [FXT: comb/setpartition-rgs-demo.cc]. Note that while the RGS correspond to a Gray code for set partitions the RGS can change in more than one position, see the left column of figure 15.2-C. About 78 million RGS per second are generated with the recursive update routine, with the iterative update the rate is about 140 million per second.

15.3.2 Max-increment RGS *

The generation of RGSs $s = [s_0, s_1, \ldots, s_{n-1}]$ where $s_k \le i + \max_{j<k}(s_j)$ is a generalization of the RGSs for set partitions (where $i = 1$). Figure 15.3-A show RGSs in lexicographic order for $i = 2$ (left) and $i = 1$ (right). The strings can be generated in lexicographic order using [FXT: `class rgs_maxincr` in comb/rgs-maxincr.h]:

```
1   class rgs_maxincr
2   {
3   public:
4       ulong *s_;  // restricted growth string
5       ulong *m_;  // m_[k-1] == max possible value for s_[k]
6       ulong n_;   // Length of strings
7       ulong i_;   // s[k] <= max_{j<k}(s[j]+i)
8       // i==1 ==> RGS for set partitions
9
10  public:
11      rgs_maxincr(ulong n, ulong i=1)
12      {
13          n_ = n;
14          m_ = new ulong[n_];
15          s_ = new ulong[n_];
16          i_ = i;
17          first();
18      }
19
20      ~rgs_maxincr()
21      {
22          delete [] m_;
23          delete [] s_;
24      }
25
26      void first()
27      {
28          ulong n = n_;
29          for (ulong k=0; k<n; ++k)  s_[k] = 0;
30          for (ulong k=0; k<n; ++k)  m_[k] = i_;
31      }
32    [--snip--]
```

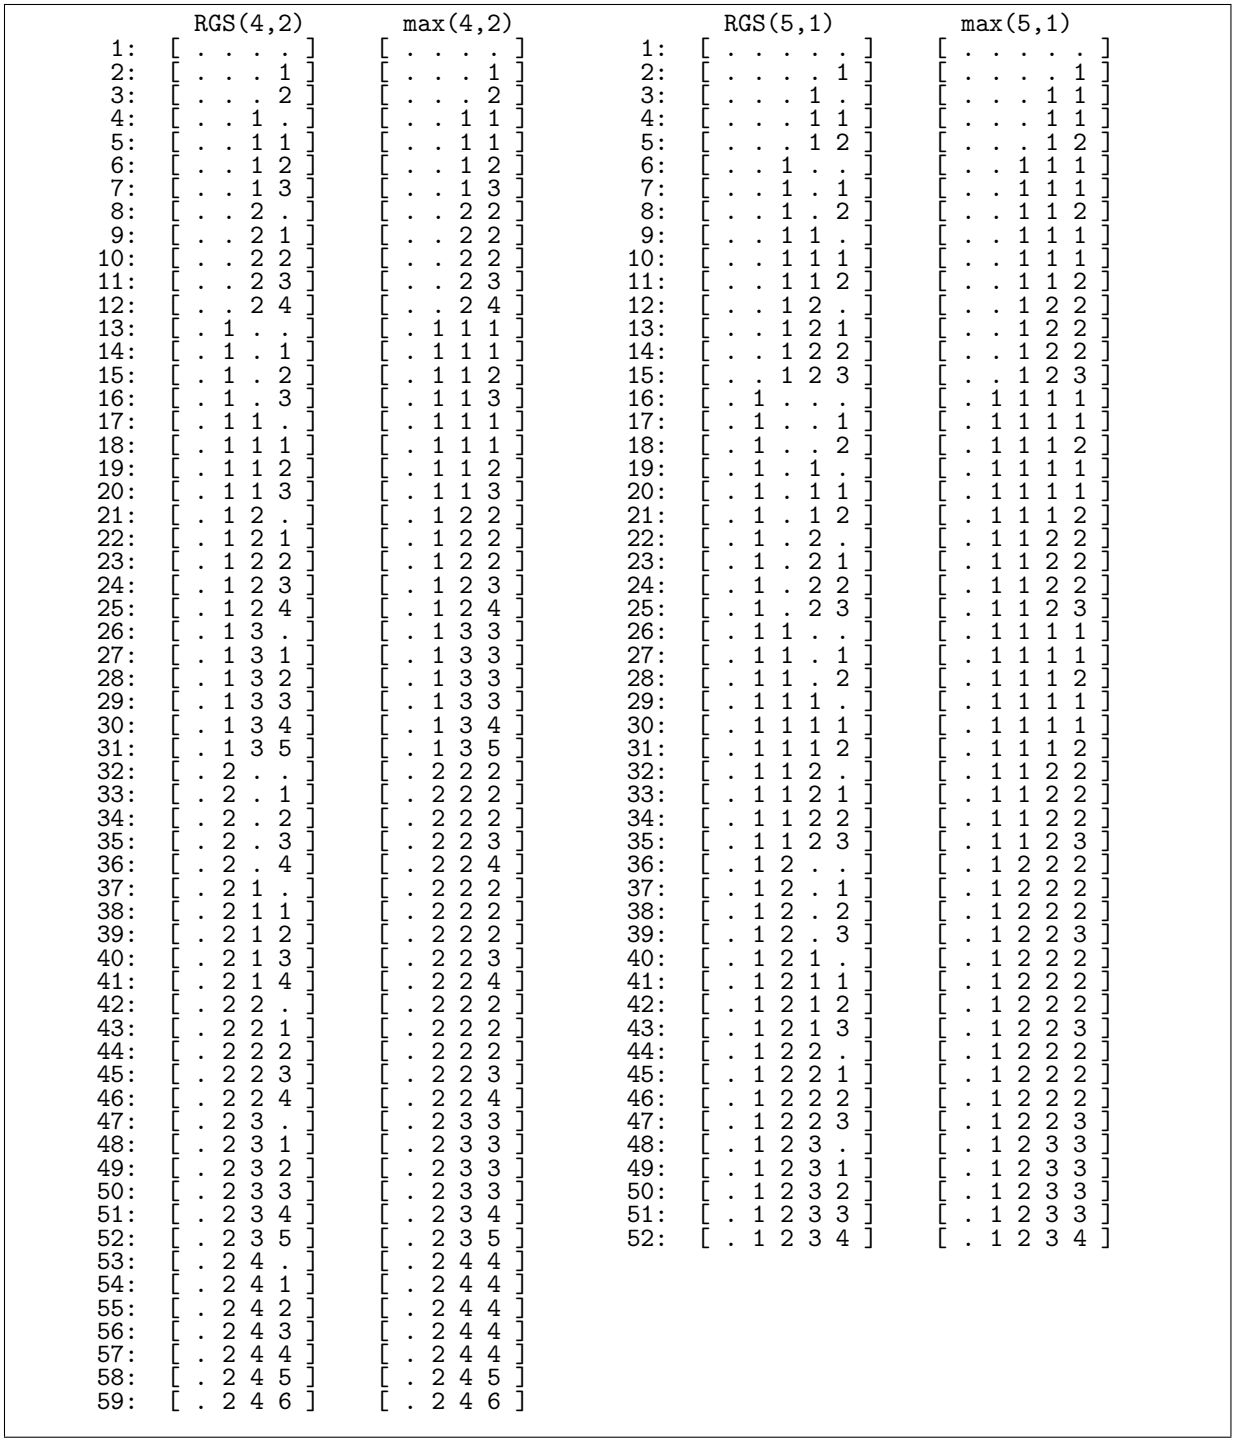

	RGS(4,2)	max(4,2)		RGS(5,1)	max(5,1)
1:	[. . . .]	[. . . .]	1:	[. . . . .]	[. . . . .]
2:	[. . . 1]	[. . . 1]	2:	[. . . . 1]	[. . . . 1]
3:	[. . . 2]	[. . . 2]	3:	[. . . 1 .]	[. . . 1 1]
4:	[. . 1 .]	[. . 1 1]	4:	[. . . 1 1]	[. . . 1 1]
5:	[. . 1 1]	[. . 1 1]	5:	[. . . 1 2]	[. . . 1 2]
6:	[. . 1 2]	[. . 1 2]	6:	[. . 1 . .]	[. . 1 1 1]
7:	[. . 1 3]	[. . 1 3]	7:	[. . 1 . 1]	[. . 1 1 1]
8:	[. . 2 .]	[. . 2 2]	8:	[. . 1 . 2]	[. . 1 1 2]
9:	[. . 2 1]	[. . 2 2]	9:	[. . 1 1 .]	[. . 1 1 1]
10:	[. . 2 2]	[. . 2 2]	10:	[. . 1 1 1]	[. . 1 1 1]
11:	[. . 2 3]	[. . 2 3]	11:	[. . 1 1 2]	[. . 1 1 2]
12:	[. . 2 4]	[. . 2 4]	12:	[. . 1 2 .]	[. . 1 2 2]
13:	[. 1 . .]	[. 1 1 1]	13:	[. . 1 2 1]	[. . 1 2 2]
14:	[. 1 . 1]	[. 1 1 1]	14:	[. . 1 2 2]	[. . 1 2 2]
15:	[. 1 . 2]	[. 1 1 2]	15:	[. . 1 2 3]	[. . 1 2 3]
16:	[. 1 . 3]	[. 1 1 3]	16:	[. 1 . . .]	[. 1 1 1 1]
17:	[. 1 1 .]	[. 1 1 1]	17:	[. 1 . . 1]	[. 1 1 1 1]
18:	[. 1 1 1]	[. 1 1 1]	18:	[. 1 . . 2]	[. 1 1 1 2]
19:	[. 1 1 2]	[. 1 1 2]	19:	[. 1 . 1 .]	[. 1 1 1 1]
20:	[. 1 1 3]	[. 1 1 3]	20:	[. 1 . 1 1]	[. 1 1 1 1]
21:	[. 1 2 .]	[. 1 2 2]	21:	[. 1 . 1 2]	[. 1 1 1 2]
22:	[. 1 2 1]	[. 1 2 2]	22:	[. 1 . 2 .]	[. 1 1 2 2]
23:	[. 1 2 2]	[. 1 2 2]	23:	[. 1 . 2 1]	[. 1 1 2 2]
24:	[. 1 2 3]	[. 1 2 3]	24:	[. 1 . 2 2]	[. 1 1 2 2]
25:	[. 1 2 4]	[. 1 2 4]	25:	[. 1 . 2 3]	[. 1 1 2 3]
26:	[. 1 3 .]	[. 1 3 3]	26:	[. 1 1 . .]	[. 1 1 1 1]
27:	[. 1 3 1]	[. 1 3 3]	27:	[. 1 1 . 1]	[. 1 1 1 1]
28:	[. 1 3 2]	[. 1 3 3]	28:	[. 1 1 . 2]	[. 1 1 1 2]
29:	[. 1 3 3]	[. 1 3 3]	29:	[. 1 1 1 .]	[. 1 1 1 1]
30:	[. 1 3 4]	[. 1 3 4]	30:	[. 1 1 1 1]	[. 1 1 1 1]
31:	[. 1 3 5]	[. 1 3 5]	31:	[. 1 1 1 2]	[. 1 1 1 2]
32:	[. 2 . .]	[. 2 2 2]	32:	[. 1 1 2 .]	[. 1 1 2 2]
33:	[. 2 . 1]	[. 2 2 2]	33:	[. 1 1 2 1]	[. 1 1 2 2]
34:	[. 2 . 2]	[. 2 2 2]	34:	[. 1 1 2 2]	[. 1 1 2 2]
35:	[. 2 . 3]	[. 2 2 3]	35:	[. 1 1 2 3]	[. 1 1 2 3]
36:	[. 2 . 4]	[. 2 2 4]	36:	[. 1 2 . .]	[. 1 2 2 2]
37:	[. 2 1 .]	[. 2 2 2]	37:	[. 1 2 . 1]	[. 1 2 2 2]
38:	[. 2 1 1]	[. 2 2 2]	38:	[. 1 2 . 2]	[. 1 2 2 2]
39:	[. 2 1 2]	[. 2 2 2]	39:	[. 1 2 . 3]	[. 1 2 2 3]
40:	[. 2 1 3]	[. 2 2 3]	40:	[. 1 2 1 .]	[. 1 2 2 2]
41:	[. 2 1 4]	[. 2 2 4]	41:	[. 1 2 1 1]	[. 1 2 2 2]
42:	[. 2 2 .]	[. 2 2 2]	42:	[. 1 2 1 2]	[. 1 2 2 2]
43:	[. 2 2 1]	[. 2 2 2]	43:	[. 1 2 1 3]	[. 1 2 2 3]
44:	[. 2 2 2]	[. 2 2 2]	44:	[. 1 2 2 .]	[. 1 2 2 2]
45:	[. 2 2 3]	[. 2 2 3]	45:	[. 1 2 2 1]	[. 1 2 2 2]
46:	[. 2 2 4]	[. 2 2 4]	46:	[. 1 2 2 2]	[. 1 2 2 2]
47:	[. 2 3 .]	[. 2 3 3]	47:	[. 1 2 2 3]	[. 1 2 2 3]
48:	[. 2 3 1]	[. 2 3 3]	48:	[. 1 2 3 .]	[. 1 2 3 3]
49:	[. 2 3 2]	[. 2 3 3]	49:	[. 1 2 3 1]	[. 1 2 3 3]
50:	[. 2 3 3]	[. 2 3 3]	50:	[. 1 2 3 2]	[. 1 2 3 3]
51:	[. 2 3 4]	[. 2 3 4]	51:	[. 1 2 3 3]	[. 1 2 3 3]
52:	[. 2 3 5]	[. 2 3 5]	52:	[. 1 2 3 4]	[. 1 2 3 4]
53:	[. 2 4 .]	[. 2 4 4]			
54:	[. 2 4 1]	[. 2 4 4]			
55:	[. 2 4 2]	[. 2 4 4]			
56:	[. 2 4 3]	[. 2 4 4]			
57:	[. 2 4 4]	[. 2 4 4]			
58:	[. 2 4 5]	[. 2 4 5]			
59:	[. 2 4 6]	[. 2 4 6]			

Figure 15.3-A: Length-4 max-increment RGS with maximal increment 2 and the corresponding array of maxima (left), and length-5 RGSs with maximal increment 1 (right). Dots denote zeros.

The computation if the successor returns the index of first (leftmost) changed element in the string. Zero is returned if the current string is the last:

```
1       ulong next()
2       {
3           ulong k = n_;
4       start:
5           --k;
6           if ( k==0 )  return 0;
7
8           ulong sk = s_[k] + 1;
9           ulong m1 = m_[k-1];
10          if ( sk > m1+i_ )  // "carry"
11          {
12              s_[k] = 0;
13              goto start;
14          }
15
16          s_[k] = sk;
17          if ( sk>m1 )  m1 = sk;
18          for (ulong j=k; j<n_; ++j )  m_[j] = m1;
19
20          return k;
21      }
22    [--snip--]
```

About 115 million RGSs per second are generated with the routine. Figure 15.3-A was created with the program [FXT: comb/rgs-maxincr-demo.cc]. The sequence of numbers of Max-increment RGSs with increment i =1, 2, 3, and 4, start

n:	0	1	2	3	4	5	6	7	8	9
i=1:	1	2	5	15	52	203	877	4140	21147	115975
i=2:	1	3	12	59	339	2210	16033	127643	1103372	10269643
i=3:	1	4	22	150	1200	10922	110844	1236326	14990380	195895202
i=4:	1	5	35	305	3125	36479	475295	6811205	106170245	1784531879

The sequence for $i = 2$ is entry A080337 of [240], it has the exponential generating function (EGF)

$$\sum_{n=0}^{\infty} B_{n,2}\,\frac{x^n}{n!} \;=\; \exp\left(x + \exp(x) + \frac{\exp(2\,x)}{2} - \frac{3}{2}\right) \tag{15.3-1}$$

The sequence of numbers of increment-3 RGSs has the EGF

$$\sum_{n=0}^{\infty} B_{n,3}\,\frac{x^n}{n!} \;=\; \exp\left(x + \exp(x) + \frac{\exp(2\,x)}{2} + \frac{\exp(3\,x)}{3} - \frac{11}{6}\right) \tag{15.3-2}$$

Omitting the empty string, we restate the EGF for the Bell numbers as

$$\sum_{n=0}^{\infty} B_{n,1}\,\frac{x^n}{n!} \;=\; \exp\left(x + \exp(x) - 1\right) \;=\; \frac{1}{0!} + \frac{2}{1!}\,x + \frac{5}{2!}\,x^2 + \frac{15}{!3}\,x^3 + \frac{52}{4!}\,x^4 + \ldots \tag{15.3-3}$$

The EGF for the increment-i RGS is

$$\sum_{n=0}^{\infty} B_{n,i}\,\frac{x^n}{n!} \;=\; \exp\left(x + \sum_{j=1}^{i} \frac{\exp(j\,x) - 1}{j}\right) \tag{15.3-4}$$

15.3.3 F-increment RGS *

To obtain a different generalization of the RGS for set partitions we rewrite the condition $s_k \le i + \max_{j<k}(s_j)$ for the RGS considered in the previous section:

$$s_k \;\le\; M(k) + i \quad \text{where} \quad M(0) \;=\; 0 \quad \text{and} \tag{15.3-5a}$$

$$M(k+1) \;=\; \begin{cases} s_{k+1} & \text{if} \quad s_{k+1} - s_k > 0 \\ M(k) & \text{else} \end{cases} \tag{15.3-5b}$$

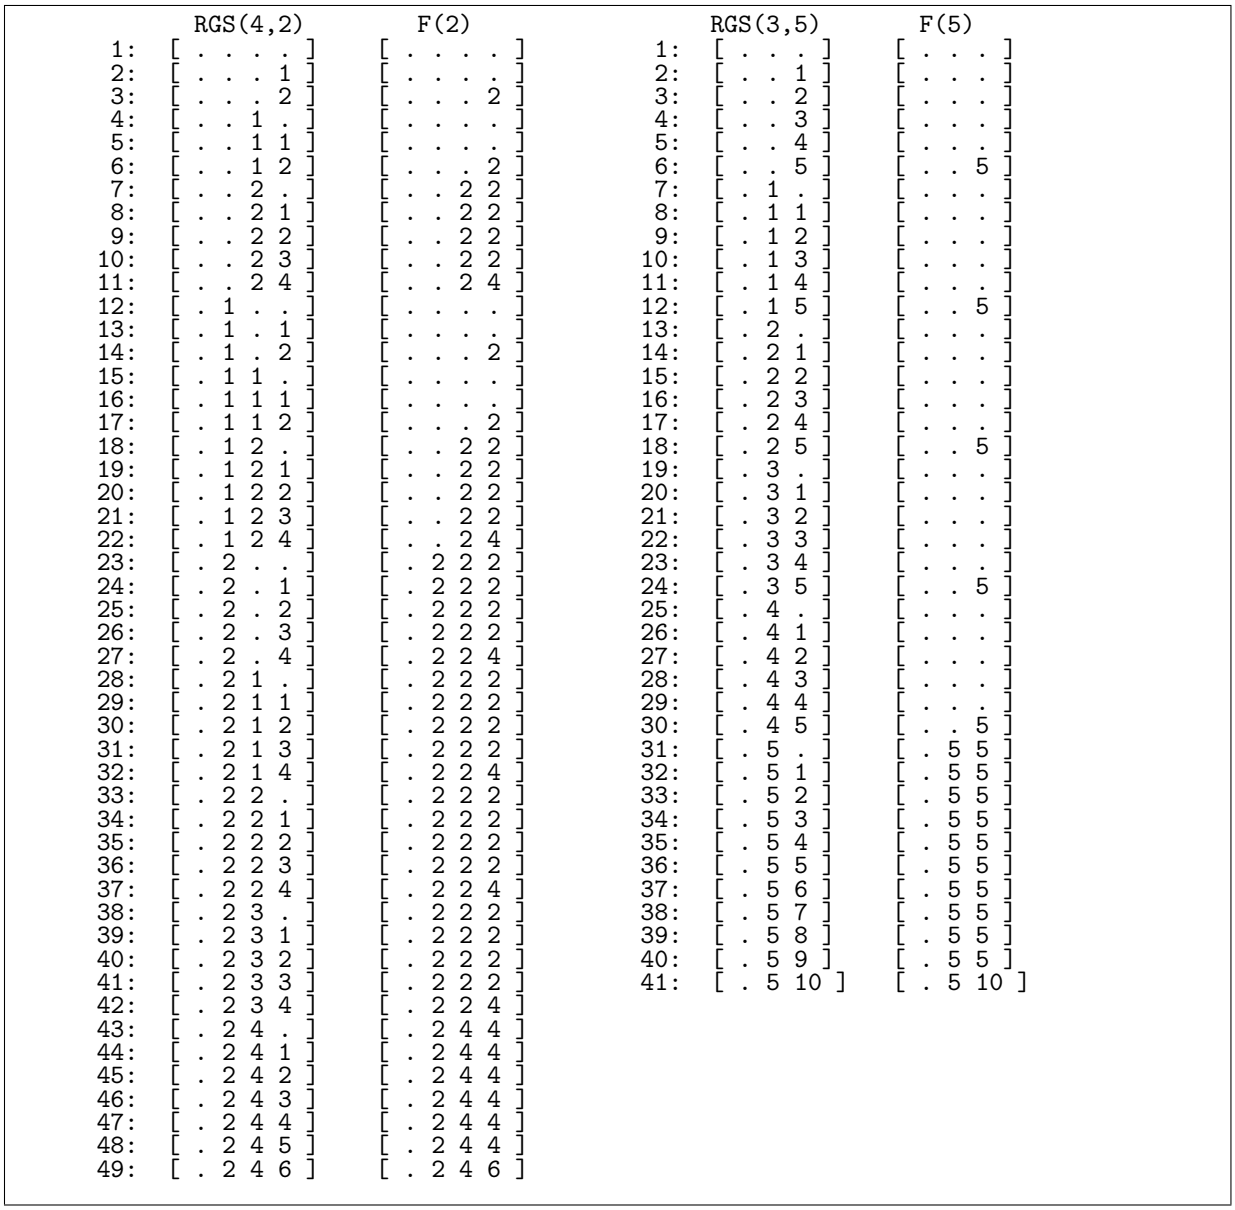

	RGS(4,2)	F(2)	RGS(3,5)	F(5)
1:	[. . . . .]	[. . . . .]	[. . . .]	[. . . .]
2:	[. . . 1]	[. . . .]	[. . 1]	[. . .]
3:	[. . . 2]	[. . . 2]	[. . 2]	[. . .]
4:	[. . 1 .]	[. . . .]	[. . 3]	[. . .]
5:	[. . 1 1]	[. . . .]	[. . 4]	[. . .]
6:	[. . 1 2]	[. . . 2]	[. . 5]	[. . 5]
7:	[. . 2 .]	[. . 2 2]	[. 1 .]	[. . .]
8:	[. . 2 1]	[. . 2 2]	[. 1 1]	[. . .]
9:	[. . 2 2]	[. . 2 2]	[. 1 2]	[. . .]
10:	[. . 2 3]	[. . 2 2]	[. 1 3]	[. . .]
11:	[. . 2 4]	[. . 2 4]	[. 1 4]	[. . .]
12:	[. 1 . .]	[. . . .]	[. 1 5]	[. . 5]
13:	[. 1 . 1]	[. . . .]	[. 2 .]	[. . .]
14:	[. 1 . 2]	[. . . 2]	[. 2 1]	[. . .]
15:	[. 1 1 .]	[. . . .]	[. 2 2]	[. . .]
16:	[. 1 1 1]	[. . . .]	[. 2 3]	[. . .]
17:	[. 1 1 2]	[. . . 2]	[. 2 4]	[. . .]
18:	[. 1 2 .]	[. . 2 2]	[. 2 5]	[. . 5]
19:	[. 1 2 1]	[. . 2 2]	[. 3 .]	[. . .]
20:	[. 1 2 2]	[. . 2 2]	[. 3 1]	[. . .]
21:	[. 1 2 3]	[. . 2 2]	[. 3 2]	[. . .]
22:	[. 1 2 4]	[. . 2 4]	[. 3 3]	[. . .]
23:	[. 2 . .]	[. 2 2 2]	[. 3 4]	[. . .]
24:	[. 2 . 1]	[. 2 2 2]	[. 3 5]	[. . 5]
25:	[. 2 . 2]	[. 2 2 2]	[. 4 .]	[. . .]
26:	[. 2 . 3]	[. 2 2 2]	[. 4 1]	[. . .]
27:	[. 2 . 4]	[. 2 2 4]	[. 4 2]	[. . .]
28:	[. 2 1 .]	[. 2 2 2]	[. 4 3]	[. . .]
29:	[. 2 1 1]	[. 2 2 2]	[. 4 4]	[. . .]
30:	[. 2 1 2]	[. 2 2 2]	[. 4 5]	[. . 5]
31:	[. 2 1 3]	[. 2 2 2]	[. 5 .]	[. 5 5]
32:	[. 2 1 4]	[. 2 2 4]	[. 5 1]	[. 5 5]
33:	[. 2 2 .]	[. 2 2 2]	[. 5 2]	[. 5 5]
34:	[. 2 2 1]	[. 2 2 2]	[. 5 3]	[. 5 5]
35:	[. 2 2 2]	[. 2 2 2]	[. 5 4]	[. 5 5]
36:	[. 2 2 3]	[. 2 2 2]	[. 5 5]	[. 5 5]
37:	[. 2 2 4]	[. 2 2 4]	[. 5 6]	[. 5 5]
38:	[. 2 3 .]	[. 2 2 2]	[. 5 7]	[. 5 5]
39:	[. 2 3 1]	[. 2 2 2]	[. 5 8]	[. 5 5]
40:	[. 2 3 2]	[. 2 2 2]	[. 5 9]	[. 5 5]
41:	[. 2 3 3]	[. 2 2 2]	[. 5 10]	[. 5 10]
42:	[. 2 3 4]	[. 2 2 4]		
43:	[. 2 4 .]	[. 2 4 4]		
44:	[. 2 4 1]	[. 2 4 4]		
45:	[. 2 4 2]	[. 2 4 4]		
46:	[. 2 4 3]	[. 2 4 4]		
47:	[. 2 4 4]	[. 2 4 4]		
48:	[. 2 4 5]	[. 2 4 4]		
49:	[. 2 4 6]	[. 2 4 6]		

Figure 15.3-B: Length-4 F-increment restricted growth strings with maximal increment 2 and the corresponding array of values of F (left), and length-3 RGSs with maximal increment 5 (right). Dots denote zeros.

The function $M(k)$ is $\max_{j<k}(s_j)$ in notational disguise. We define *F-increment* RGSs with respect to a function F as follows:

$$s_k \;\le\; F(k)+i \quad \text{where} \quad F(0) \;=\; 0 \quad \text{and} \tag{15.3-6a}$$

$$F(k+1) \;=\; \begin{cases} s_{k+1} & \text{if} \quad s_{k+1}-s_k = i \\ F(k) & \text{else} \end{cases} \tag{15.3-6b}$$

The function $F(k)$ is a 'maximum' that is increased only if the last increase $(s_k - s_{k-1})$ was maximal. For $i = 1$ we obtain the RGSs for set partitions. Figure 15.3-B shows all length-4 F-increment RGSs for $i = 2$ (left), and all length-3 RGSs for $i = 5$ (right), together with the arrays of F-values. The listings were created with the program [FXT: comb/rgs-fincr-demo.cc]. It uses the implementation [FXT: `class rgs_fincr` in comb/rgs-fincr.h]:

```
1   class rgs_fincr
2   {
3   public:
4       ulong *s_;  // restricted growth string
5       ulong *f_;  // values F(k)
6       ulong n_;   // Length of strings
7       ulong i_;   // s[k] <= f[k]+i
8     [--snip--]
9
10      ulong next()
11      // Return index of first changed element in s[],
12      // Return zero if current string is the last
13      {
14          ulong k = n_;
15
16      start:
17          --k;
18          if ( k==0 )  return 0;
19
20          ulong sk = s_[k] + 1;
21          ulong m1 = f_[k-1];
22          ulong mp = m1 + i_;
23          if ( sk > mp )  // "carry"
24          {
25              s_[k] = 0;
26              goto start;
27          }
28
29          s_[k] = sk;
30          if ( sk==mp )  m1 += i_;
31          for (ulong j=k; j<n_; ++j )  f_[j] = m1;
32
33          return k;
34      }
35    [--snip--]
```

The sequences of numbers of F-increment RGSs with increments i =1, 2, 3, and 4, start

```
  n:  0  1   2    3     4     5      6       7        8         9
i=1:  1  2   5   15    52   203    877    4140    21147    115975
i=2:  1  3  11   49   257  1539  10299   75905   609441   5284451
i=3:  1  4  19  109   742  5815  51193  498118  5296321  60987817
i=4:  1  5  29  201  1657 15821 170389 2032785 26546673 376085653
```

These are respectively entries A000110 (Bell numbers), A004211, A004212, and A004213 of [240]. In general, the number $F_{n,i}$ of F-increment RGSs (length n, with increment i) is

$$F_{n,i} \;=\; \sum_{k=0}^{n} i^{n-k}\, S(n,k) \tag{15.3-7}$$

where $S(n,k)$ are the Stirling numbers of the second kind. The exponential generating functions are

$$\sum_{n=0}^{\infty} F_{n,i}\, \frac{x^n}{n!} \;=\; \exp\left(\frac{\exp(i\,x)-1}{i}\right) \tag{15.3-8}$$

The ordinary generating functions are

$$\sum_{n=0}^{\infty} F_{n,i}\, x^n \quad = \quad \sum_{n=0}^{\infty} \frac{x^n}{\prod_{k=1}^{n} 1 - i\,k\,x} \tag{15.3-9}$$

15.3.4 K-increment RGS *

```
 1:  [ . . . . ]
 2:  [ . . . 1 ]    11:  [ . . 2 1 ]    20:  [ . 1 1 . ]    29:  [ . 1 2 4 ]
 3:  [ . . . 2 ]    12:  [ . . 2 2 ]    21:  [ . 1 1 1 ]    30:  [ . 1 2 5 ]
 4:  [ . . . 3 ]    13:  [ . . 2 3 ]    22:  [ . 1 1 2 ]    31:  [ . 1 3 . ]
 5:  [ . . 1 . ]    14:  [ . . 2 4 ]    23:  [ . 1 1 3 ]    32:  [ . 1 3 1 ]
 6:  [ . . 1 1 ]    15:  [ . . 2 5 ]    24:  [ . 1 1 4 ]    33:  [ . 1 3 2 ]
 7:  [ . . 1 2 ]    16:  [ . 1 . . ]    25:  [ . 1 2 . ]    34:  [ . 1 3 3 ]
 8:  [ . . 1 3 ]    17:  [ . 1 . 1 ]    26:  [ . 1 2 1 ]    35:  [ . 1 3 4 ]
 9:  [ . . 1 4 ]    18:  [ . 1 . 2 ]    27:  [ . 1 2 2 ]    36:  [ . 1 3 5 ]
10:  [ . . 2 . ]    19:  [ . 1 . 3 ]    28:  [ . 1 2 3 ]    37:  [ . 1 3 6 ]
```

Figure 15.3-C: The 37 K-increment RGS of length 4 in lexicographic order.

We mention yet another type of restricted growth strings, the *K-increment* RGS, which satisfy

$$s_k \quad \leq \quad s_{k-1} + k \tag{15.3-10}$$

An implementation for their generation in lexicographic order is given in [FXT: comb/rgs-kincr.h]:

```
class rgs_kincr
{
public:
    ulong *s_;  // restricted growth string
    ulong n_;   // Length of strings
  [--snip--]

    ulong next()
    // Return index of first changed element in s[],
    // Return zero if current string is the last
    {
        ulong k = n_;

    start:
        --k;
        if ( k==0 )  return 0;

        ulong sk = s_[k] + 1;
        ulong mp = s_[k-1] + k;
        if ( sk > mp )  // "carry"
        {
            s_[k] = 0;
            goto start;
        }

        s_[k] = sk;
        return k;
    }
  [--snip--]
```

The sequence of the numbers of K-increment RGS of length n is entry A107877 of [240]:

```
n:  0  1  2  3   4   5    6     7      8       9        10
    1  1  2  7  37  268  2496  28612  391189  6230646  113521387
```

The strings of length 4 are shown in figure 15.3-C. They can be generated with the program [FXT: comb/rgs-kincr-demo.cc].

Chapter 16

A string substitution engine

```
Number of symbols = 3
Start: x
Rules:
  x --> A
  A --> Bx
  B --> Bx
-------------
0:   (#=1)
  x
1:   (#=1)
  A
2:   (#=2)
  Bx
3:   (#=3)
  BxA
4:   (#=5)
  BxABx
5:   (#=8)
  BxABxBxA
6:   (#=13)
  BxABxBxABxABx
7:   (#=21)
  BxABxBxABxABxBxABxBxA
8:   (#=34)
  BxABxBxABxABxBxABxBxABxABxBxABxABx
9:   (#=55)
  BxABxBxABxABxBxABxBxABxABxBxABxABxBxABxBxABxABxBxABxBxA
```

Figure 16.0-A: Rules, axiom, and first steps of the evolution for a specific substitution engine.

String substitution: A finite set of symbols ('alphabet'), a set of substitution rules that map symbols to words (strings of symbols) and a start value ('axiom'). The rules have to be applied in parallel. An example, let $\{x, A, B\}$ be the symbols and $\{x \mapsto A, A \mapsto Bx, B \mapsto Bx\}$ the substitution rules. Start with a single x. The observed evolution is shown in figure 16.0-A. It is trivial to implement such a specific system, for the above:

```
1   void fx(ulong n);
2   void fA(ulong n);
3   void fB(ulong n);
4
5   void fx(ulong n)
6   {
7       if ( 0==n ) { cout << "x"; return; }
8       fA(n-1);
9   }
10
11  void fA(ulong n)
12  {
13      if ( 0==n ) { cout << "A"; return; }
14      fB(n-1);  fx(n-1);
15  }
16
17  void fB(ulong n)
```

```
{
    if ( 0==n ) { cout << "B"; return; }
    fB(n-1);  fx(n-1);
}

int main()
{
    ulong n = 10;  // max depth

    for (ulong k=0; k<=n; ++k)
    {
        fx(k);
        cout << endl;
    }

    return 0;
}
```

A utility class to create string substitution engines at run time is given in [FXT: `class string_subst` in comb/stringsubst.h]:

```
class string_subst
{
public:
    // example values generate rabbit sequence:
    ulong nsym_;    // # of symbols
    // == 2
    char *symbol_;   // alphabet
    // == { '0', '1' };
    char **symrule_; // symrule_[i] string to replace i-th symbol with
    // == { "0", "1",  "1", "10" };  for  0 |-> 1,  1 |-> 10

    ulong xlate_[256];  // translate char --> rule
    //  'a' --> symrule[xlate_['a']]

    ulong cmax_; // max string length
    char  *cc_;  // string to hold result
    ulong ctc_;  // count chars of result actually produced


public:
    string_subst(int cmax, int nsym, char*const* symrule);
```

Rules and symbols have to be supplied on construction. The member function `subst` takes an axiom and the number of generations as arguments and returns the numbers of produced characters. It calls the recursive function `do_subst`:

```
    void do_subst(ulong n, const char *rule)
    {
        if ( 0==n )  // add symbols to string:
        {
            for (ulong i=0;  ctc_<cmax_;  ++i)
            {
                char c = rule[i];
                if ( 0==c )  break;
                cc_[ctc_] = c;
                ++ctc_;
            }
        }
        else  // recurse:
        {
            for (ulong i=0;  0!=rule[i];  ++i)
            {
                ulong r = (ulong)rule[i];
                ulong x = xlate_[r];
                do_subst(n-1, symrule_[x]);  // recursion
            }
        }
    }

    ulong subst(ulong maxn, const char *start)
    // maxn:=number of generations,  start:=axiom
    {
        ctc_ = 0;
        do_subst(maxn, start);
        cc_[ctc_] = 0;  // terminate string
        return  ctc_;
```

```
31      }
```

The resulting string is stored in an array of characters that can be accessed via the member function `string()`. The characters that do not fit into the array are silently discarded. Using an array of one million characters, the rules and axiom as follows

```
Number of symbols = 5
Start: a
Rules:
  a --> -bF+aFa+Fb-
  b --> +aF-bFb-Fa+
  + --> +
  - --> -
  F --> F
```

Then one obtains (displaying only the string lengths from the fourth generation on):

```
 0:   (#=1)
   a
 1:   (#=11)
   -bF+aFa+Fb-
 2:   (#=51)
   -+aF-bFb-Fa+F+-bF+aFa+Fb-F-bF+aFa+Fb-+F+aF-bFb-Fa+-
 3:   (#=211)
   -+-bF+aFa+Fb-F-+aF-bFb-Fa+F+aF-bFb-Fa+-F-bF+aFa+Fb-+F+-+aF-bFb-  \
    Fa+F+-bF+aFa+Fb-F-bF+aFa+Fb-+F+aF-bFb-Fa+-F-+aF-bFb-Fa+F+-bF+a  \
    Fa+Fb-F-bF+aFa+Fb-+F+aF-bFb-Fa+-+F+-bF+aFa+Fb-F-+aF-bFb-Fa+F+a  \
    F-bFb-Fa+-F-bF+aFa+Fb-+-
 4:   (#=851)
 5:   (#=3411)
 6:   (#=13651)
 7:   (#=54611)
 8:   (#=218451)
 9:   (#=873811)
10:   (#=1000000)  // <-- truncation
11:   (#=1000000)
```

The computation takes a split second only. The strings generated by the shown rule correspond to the Hilbert curves in successively finer resolutions if one initializes

```
position:         (x, y)   := (0,0)
direction:        (dx, dy) := (1,0)
```

and identifies:

```
'-'  ==  "turn left":   { t= dx; dx=-dy; dy=t; }
'+'  ==  "turn right":  { t=-dx; dx= dy; dy=t; }
'F'  ==  "advance":     { line(x, y, x+dx, y+dy);  x+=dx; y+=dy; }
either 'a' or 'b'  ==   (ignore)
```

The program [FXT: ds/stringsubst-demo.cc] prints evolutions for several rules. A 3-dimensional Hilbert (space-filling) curve is generated by the rule `a` $\mapsto$ `^<aF^<aFa-F^>>aFavF+>>aFa-F>a->`, all other symbols are mapped to themselves. The curve is obtained by interpreting the symbols (unequal to `F`) of the string as rotations [FXT: ds/stringsubst-hilbert3d-demo.cc].

The number of symbols that occur in each generation can be computed by a simple matrix power. With i different symbols let $n_j(k)$ be the number of occurrences of the j-th symbol at generation k. Then $n_j(0)$ is determined by the axiom and

$$\vec{N}(k) \quad := \quad [n_1(k), n_2(k), \ldots, n_i(k)]^T \tag{16.0-1}$$

can be computed as

$$\vec{N}(k) \quad = \quad M^k \, \vec{N}(0) \tag{16.0-2}$$

where M is a matrix with columns according to the occurrences of the symbols in the transfer rules. For example, with the rules and axiom

```
Number of symbols = 3
Start: x
Rules:                  x  A  B
  x --> A    1 --> [0, 1, 0]
  A --> Bx   2 --> [1, 0, 1]  =: transpose(M)
  B --> Bx   3 --> [1, 0, 1]
```

we obtain

$$M = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix} \qquad M^k = \begin{bmatrix} F_{k-1} & F_k & F_k \\ F_{k-2} & F_{k-1} & F_{k-1} \\ F_{k-1} & F_k & F_k \end{bmatrix} \qquad k \geq 2 \tag{16.0-3}$$

where F_k is the k-th Fibonacci number. For the tenth generation, $k = 10$:

$$M^{10}\,[1,0,0]^T = \begin{bmatrix} 34 & 55 & 55 \\ 21 & 34 & 34 \\ 34 & 55 & 55 \end{bmatrix} [1,0,1]^T = [34,21,34]^T \tag{16.0-4}$$

So we have 34 '`x`', 21 '`A`' and 34 '`B`' or a total of 89 symbols.

```
Number of symbols = 5
Start: a
Rules:
  a --> -bF+aFa+Fb-   1 --> [2, 2, 2, 2, 3]
  b --> +aF-bFb-Fa+   2 --> [2, 2, 2, 2, 3]
  + --> +             3 --> [0, 0, 1, 0, 0] =: transpose(m)
  - --> -             4 --> [0, 0, 0, 1, 0]
  F --> F             5 --> [0, 0, 0, 0, 1]
```

k	total	[n(a),	n(b),	n(+),	n(-),	n(F)]
0	1	[1,	0,	0,	0,	0]
1	11	[2,	2,	2,	2,	3]
2	51	[8,	8,	10,	10,	15]
3	211	[32,	32,	42,	42,	63]
4	851	[128,	128,	170,	170,	255]
5	3411	[512,	512,	682,	682,	1023]
6	13651	[2048,	2048,	2730,	2730,	4095]
7	54611	[8192,	8192,	10922,	10922,	16383]
8	218451	[32768,	32768,	43690,	43690,	65535]
9	873811	[131072,	131072,	174762,	174762,	262143]
10	3495251	[524288,	524288,	699050,	699050,	1048575]

Figure 16.0-B: Rules for a string substitution engine for a 3-dimensional Hilbert curve (top), and number of occurrences of the symbols after k iterations (bottom).

The rules for the engine corresponding to a 3D Hilbert curve are given in figure 16.0-B (top), the number of symbols after the k-th iteration are given at the bottom of the figure. In chapter 36 on page 713 the string substitution idea is used to construct fast iterations for various functions whose power series are determined by the limiting string.

An efficient algorithm for finding the n-th letter of the string produced by k iterations of a substitution rule is described in [237]. The problem of the identification of a finite automaton (substitution engine) that generates a given infinite string is discussed in [178]. The mathematical properties of sequences generated by finite automata are studied in [10].

Chapter 17

Necklaces and Lyndon words

A sequence that is minimal among all its cyclic rotations is called a *necklace* (see section 3.8.1 on page 136 for the definition in terms of equivalence classes). When there are k possible values for each element one talks about an n-bead, k-color (or k-ary length-n) necklaces. We restrict our attention to the case were only two sorts of beads are allowed and represent them by 0 and 1.

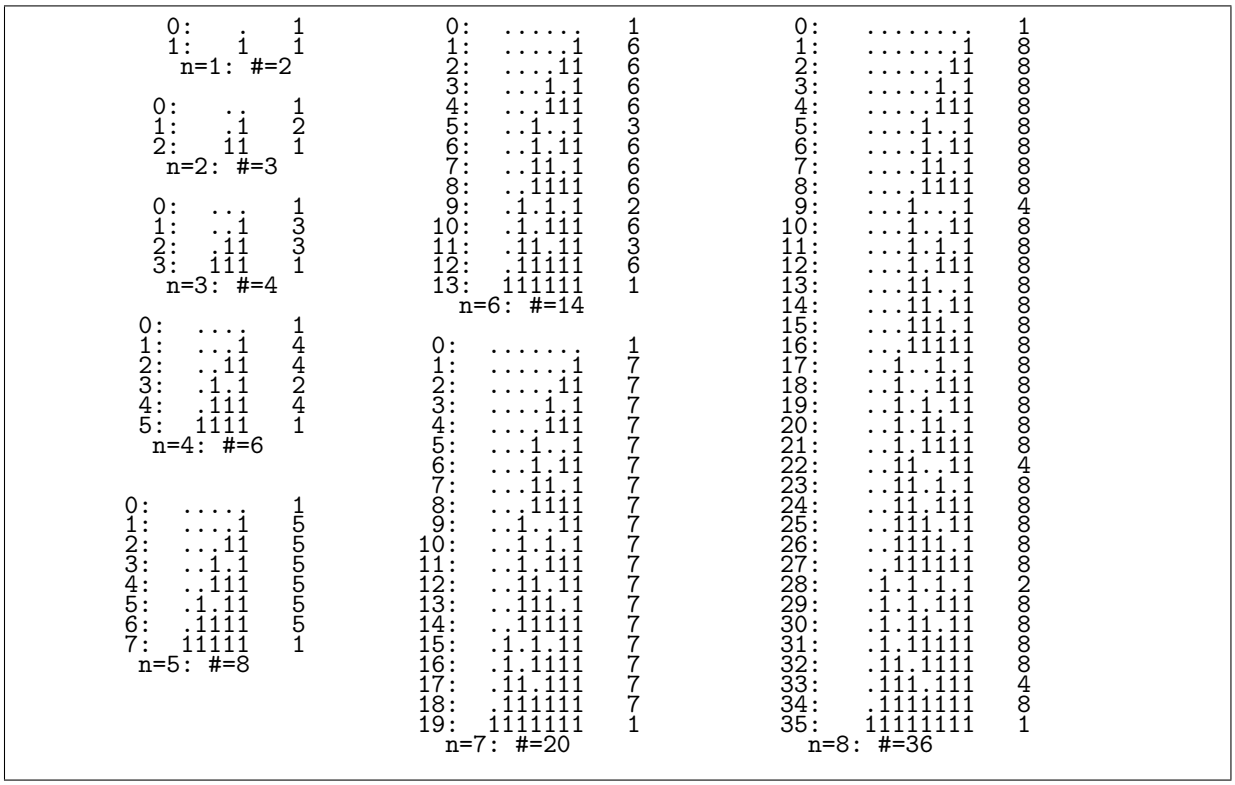

Figure 17.0-A: All binary necklaces of lengths up to 8 and their periods. Dots represent zeros.

Scanning all binary words of length n as to whether they are necklaces can easily be achieved by testing whether the word `x` is equal to the return value of the function `bit_cyclic_min(x,n)` shown in section 1.13 on page 29. For n up to 8 one obtains the sequences of binary necklaces shown in figure 17.0-A. As 2^n words have to be tested this approach gets ineffective for large n. Luckily there is both a much better algorithm for generating all necklaces and a formula for their number.

Not all necklaces are created equal. Each necklace can be assigned a period that is a divisor of the length.

That period is the smallest (nonzero) cyclic shift that transforms the word into itself. The periods are given directly right to each necklace in figure 17.0-A. For n prime the only periodic necklaces are those two that contain all ones or zeros. Aperiodic (or equivalently, period equals length) necklaces are called *Lyndon words*.

For a length-n binary word x the function `bit_cyclic_period(x,n)` from section 1.13 on page 29 returns the period of the word.

17.1 Generating all necklaces

We give several methods to generate all necklaces of a given size. An efficient algorithm for the generation of bracelets (see section 3.8.1.4 on page 137) is given in [231].

17.1.1 The FKM algorithm

```
 1:    [ . . . . ]   j=1  N
 2:    [ . . . 1 ]   j=4  N  L
 3:    [ . . . 2 ]   j=4  N  L
 4:    [ . . 1 . ]   j=3
 5:    [ . . 1 1 ]   j=4  N  L
 6:    [ . . 1 2 ]   j=4  N  L
 7:    [ . . 2 . ]   j=3
 8:    [ . . 2 1 ]   j=4  N  L
 9:    [ . . 2 2 ]   j=4  N  L
10:    [ . 1 . 1 ]   j=2  N
11:    [ . 1 . 2 ]   j=4  N  L
12:    [ . 1 1 . ]   j=3
13:    [ . 1 1 1 ]   j=4  N  L
14:    [ . 1 1 2 ]   j=4  N  L
15:    [ . 1 2 . ]   j=3
16:    [ . 1 2 1 ]   j=4  N  L
17:    [ . 1 2 2 ]   j=4  N  L
18:    [ . 2 . 2 ]   j=2  N
19:    [ . 2 1 . ]   j=3
20:    [ . 2 1 1 ]   j=4  N  L
21:    [ . 2 1 2 ]   j=4  N  L
22:    [ . 2 2 . ]   j=3
23:    [ . 2 2 1 ]   j=4  N  L
24:    [ . 2 2 2 ]   j=4  N  L
25:    [ 1 1 1 1 ]   j=1  N
26:    [ 1 1 1 2 ]   j=4  N  L
27:    [ 1 1 2 1 ]   j=3
28:    [ 1 1 2 2 ]   j=4  N  L
29:    [ 1 2 1 2 ]   j=2  N
30:    [ 1 2 2 1 ]   j=3
31:    [ 1 2 2 2 ]   j=4  N  L
32:    [ 2 2 2 2 ]   j=1  N
32 (4, 3) pre-necklaces.
24 necklaces and 18 Lyndon words.
```

```
 1:    [ . . . . . . ]   j=1  N
 2:    [ . . . . . 1 ]   j=6  N  L
 3:    [ . . . . 1 . ]   j=5
 4:    [ . . . . 1 1 ]   j=6  N  L
 5:    [ . . . 1 . . ]   j=4
 6:    [ . . . 1 . 1 ]   j=6  N  L
 7:    [ . . . 1 1 . ]   j=5
 8:    [ . . . 1 1 1 ]   j=6  N  L
 9:    [ . . 1 . . 1 ]   j=3  N
10:    [ . . 1 . 1 . ]   j=5
11:    [ . . 1 . 1 1 ]   j=6  N  L
12:    [ . . 1 1 . . ]   j=4
13:    [ . . 1 1 . 1 ]   j=6  N  L
14:    [ . . 1 1 1 . ]   j=5
15:    [ . . 1 1 1 1 ]   j=6  N  L
16:    [ . 1 . 1 . 1 ]   j=2  N
17:    [ . 1 . 1 1 . ]   j=5
18:    [ . 1 . 1 1 1 ]   j=6  N  L
19:    [ . 1 1 . 1 1 ]   j=3  N
20:    [ . 1 1 1 . 1 ]   j=4
21:    [ . 1 1 1 1 . ]   j=5
22:    [ . 1 1 1 1 1 ]   j=6  N  L
23:    [ 1 1 1 1 1 1 ]   j=1  N
23 (6, 2) pre-necklaces.
14 necklaces and 9 Lyndon words.
```

Figure 17.1-A: Ternary length-4 (left) and binary length-6 (right) pre-necklaces as generated by the FKM algorithm. Dots are used for zeros, necklaces are marked with '`N`', Lyndon words with '`L`'.

The following algorithm for generation of all necklaces actually produces *pre-necklaces* a subset of which are the necklaces. A pre-necklace is a string that is the prefix of some necklace. The *FKM algorithm* (for Fredericksen, Kessler, Maiorana) to generate all k-ary length-n pre-necklaces proceeds as follows:

1. Initialize the word $F = [f_1, f_2, \ldots, f_n]$ to all zeros. Set $j = 1$.
2. (Visit pre-necklace F. If j divides n then F is a necklace. If j equals n then F is a Lyndon word.)
3. Find the largest index j so that $f_j < k-1$. If there is no such index (then $F = [k-1, k-1, \ldots, k-1]$, the last necklace), then terminate.
4. Increment f_j. Fill the suffix starting at f_{j+1} with copies of $[f_1, \ldots, f_j]$. Goto step 2.

The crucial steps are [FXT: comb/necklace-fkm-demo.cc]:

```
1       for (ulong i=1; i<=n; ++i)  f[i] = 0;  // Initialize to zero
2       bool nq = 1;  // whether pre-necklace is a necklace
3       bool lq = 0;  // whether pre-necklace is a Lyndon word
4       ulong j = 1;
5       while ( 1 )
6       {
7           // Print necklace:
8           cout << setw(4) << pct << ":";
9           print_vec("    ", f+1, n, true);
10          cout << "   j=" << j;
11          if ( nq )  cout << "  N";
12          if ( lq )  cout << "  L";
13          cout << endl;
14
15          // Find largest index where we can increment:
16          j = n;
17          while ( f[j]==k-1 )  { --j; };
18
19          if ( j==0 )  break;
20
21          ++f[j];
22
23          // Copy periodically:
24          for (ulong i=1,t=j+1; t<=n; ++i,++t)  f[t] = f[i];
25
26          nq = ( (n%j)==0 );  // necklace if j divides n
27          lq = ( j==n );      // Lyndon word if j equals n
28      }
```

Two example runs are shown in figure 17.1-A. An efficient implementation of the algorithm is [FXT: `class necklace` in comb/necklace.h]:

```
1   class necklace
2   {
3   public:
4       ulong *a_;  // the string
5       ulong *dv_; // delta sequence of divisors of n
6       ulong n_;   // length of strings
7       ulong m1_;  // m-ary strings, m1=m-1
8       ulong j_;   // period of the word (if necklaces)
9
10  public:
11      necklace(ulong m, ulong n)
12      {
13          n_ = ( n ? n : 1 );  // at least one
14          m1_ = ( m>1 ? m-1 : 1); // at least two
15          a_ = new ulong[n_+1];
16          dv_ = new ulong[n_+1];
17          for (ulong j=1; j<=n; ++j)  dv_[j] = ( 0==(n_%j ) );  // divisors
18          first();
19      }
20    [--snip--]
21
22      void first()
23      {
24          for (ulong j=0; j<=n_; ++j)  a_[j] = 0;
25          j_ = 1;
26      }
27    [--snip--]
```

The method to compute the next pre-necklace is

```
1       ulong next_pre()  // next pre-necklace
2       // return j (zero when finished)
```

```
3       {
4           // Find rightmost digit that can be incremented:
5           ulong j = n_;
6           while ( a_[j] == m1_ )  { --j; }
7
8           // Increment:
9           // if ( 0==j_ )   return 0;  // last
10          ++a_[j];
11
12          // Copy periodically:
13          for (ulong k=j+1; k<=n_; ++k)  a_[k] = a_[k-j];
14
15          j_ = j;
16          return  j;
17      }
```

Note the commented out return with the last word, this gives a speedup (and no harm is done with the following copying). The array `dv` allows determination whether the current pre-necklace is also a necklace (or Lyndon word) via simple lookups:

```
1       bool is_necklace()  const
2       {
3           return ( 0!=dv_[j_] );  // whether j divides n
4       }
5
6       bool is_lyn()  const
7       {
8           return ( j_==n_ );  // whether j equals n
9       }
10
```

Thereby we have methods for the computation of the next necklace or Lyndon word:

```
1       ulong next()  // next necklace
2       {
3           do
4           {
5               next_pre();
6               if ( 0==j_ )  return 0;
7           }
8           while ( 0==dv_[j_] );  // until j divides n
9           return j_;
10      }
11
12      ulong next_lyn()  // next Lyndon word
13      {
14          do
15          {
16              next_pre();
17              if ( 0==j_ )  return 0;
18          }
19          while ( j_==n_ );  // until j equals n
20          return j_;  // == n
21      }
22  };
```

The rate of generation for pre-necklaces is about 98 M/s for base 2, 140 M/s for base 3, and 180 M/s for base 4 [FXT: comb/necklace-demo.cc]. An specialization of the algorithm for binary necklaces is [FXT: `class binary_necklace` in comb/binary-necklace.h]. It can generate pre-necklaces at about 124 M/s (when arrays are used instead of pointers), see [FXT: comb/binary-necklace-demo.cc]. A version of the algorithm that produces the binary necklaces as bits of a word is given in section 1.13.3 on page 31.

The binary necklaces of length n can be used as cycle leaders in the length-2^n zip permutation (and its inverse) that is discussed in section 2.5 on page 101. An algorithm for the generation of all irreducible binary polynomials via Lyndon words is described in section 38.10 on page 849.

```
 0 :   a= ......1 =    1    ==   ......1
 1 :   a= .....11 =    3    ==   .....11
 2 :   a= ...1..1 =    9    ==   ...1..1
 3 :   a= ..11.11 =   27    ==   ..11.11
 4 :   a= 1.1...1 =   81    ==   ...11.1
 5 :   a= 111.1.. =  116    ==   ..111.1
 6 :   a= 1.1111. =   94    ==   .1.1111
 7 :   a= ..111.. =   28    ==   ....111
 8 :   a= 1.1.1.. =   84    ==   ..1.1.1
 9 :   a= 11111.1 =  125    ==   .111111
10 :   a= 1111..1 =  121    ==   ..11111
11 :   a= 11.11.1 =  109    ==   .11.111
12 :   a= 1..1..1 =   73    ==   ..1..11
13 :   a= 1.111.. =   92    ==   ..1.111
14 :   a= ..1.11. =   22    ==   ...1.11
15 :   a= 1....1. =   66    ==   ....1.1
16 :   a= 1...111 =   71    ==   ...1111
17 :   a= 1.1.11. =   86    ==   .1.1.11
18 :   a= ....1.. =    4    ==   ......1 <--= sequence restarts
19 :   a= ...11.. =   12    ==   .....11
20 :   a= .1..1.. =   36    ==   ...1..1
21 :   a= 11.11.. =  108    ==   ..11.11
22 :   a= 1...11. =   70    ==   ...11.1
23 :   a= 1.1..11 =   83    ==   ..111.1
24 :   a= 1111.1. =  122    ==   .1.1111
25 :   a= 111.... =  112    ==   ....111
[--snip--]
```

Figure 17.1-B: Generation of all (18) 7-bit Lyndon words as binary representations of the powers modulo 127 of the primitive root 3. The right column gives the cyclic minima. Dots are used for zeros.

17.1.2 Binary Lyndon words with length a Mersenne exponent

The length-n binary Lyndon words for n an exponent of a Mersenne prime $M_n = 2^n - 1$ can be generated efficiently as binary expansions of the powers of a primitive root r of m until the second word with just one bit is reached. With $n = 7$, $m = 127$ and the primitive root $r = 3$ we get the sequence shown in figure 17.1-B. The sequence of minimal primitive roots r_n of the first Mersenne primes $M_n = 2^n - 1$ is entry A096393 of [240]:

```
 2: 2      17: 3        107: 3
 3: 3      19: 3        127: 43
 5: 3      31: 7        521: 3
 7: 3      61: 37       607: 5  <--= 5 is a primitive root of 2**607-1
13: 17     89: 3       1279: 5
```

17.1.3 A constant amortized time (CAT) algorithm

A constant amortized time (CAT) algorithm to generate all k-ary length-n (pre-)necklaces is given in [79]. The crucial part of a recursive algorithm [FXT: comb/necklace-cat-demo.cc] is the function

```
1   ulong K, N;  // K-ary pre-necklaces of length N
2   ulong f[N];
3   void crsms_gen(ulong n, ulong j)
4   {
5       if ( n > N )  visit(j);  // pre-necklace in f[1,...,N]
6       else
7       {
8           f[n] = f[n-j];
9           crsms_gen(n+1, j);
10
11          for (ulong i=f[n-j]+1; i<K; ++i)
12          {
13              f[n] = i;
14              crsms_gen(n+1, n);
15          }
16      }
17  }
```

After initializing the array with zeros the function must be called with both arguments equal to one. The routine generates about 71 million binary pre-necklaces per second. Ternary and 5-ary pre-necklaces are generated at a rate of about 100 and 113 million per second, respectively.

17.1.4 An order with fewer transitions

```
 1:  .......1        11:  ...11111        21:  ..1.1.11
 2:  ......11        12:  ...111.1        22:  ..1.1111
 3:  .....111        13:  ...1.1.1        23:  ..1.11.1
 4:  .....1.1        14:  ...1.111        24:  ..1..111 <<+1
 5:  ....11.1        15:  ...1..11        25:  ..1..1.1
 6:  ....1111        16:  ..11.111 <<+1   26:  .11.1111 <<+2
 7:  ....1.11        17:  ..11.1.1        27:  .1111111
 8:  ....1..1        18:  ..1111.1        28:  .1.11.11 <<+1
 9:  ...11..1        19:  ..111111        29:  .1.11111
10:  ...11.11        20:  ..111.11        30:  .1.1.111
```

Figure 17.1-C: The 30 binary 8-bit Lyndon words in an order with few changes between successive words. Transitions where more than one bit changes are marked with a '<<'.

$n:$	X_n	$n:$	X_n	$n:$	X_n	$n:$	X_n	$n:$	X_n
1:	0	7:	2	13:	95	19:	2598	25:	85449
2:	0	8:	5	14:	163	20:	4546	26:	155431
3:	0	9:	11	15:	290	21:	8135	27:	284886
4:	0	10:	15	16:	479	22:	14427	28:	522292
5:	1	11:	34	17:	859	23:	26122	29:	963237
6:	1	12:	54	18:	1450	24:	46957	30:	1778145

Figure 17.1-D: Excess (with respect to Gray code) of the number of bits changed.

The following routine generates the binary pre-necklaces words in the order that would be obtained by selecting valid words from the binary Gray code:

```
1   void xgen(ulong n, ulong j, int x=+1)
2   {
3       if ( n > N )  visit(j);
4       else
5       {
6           if ( -1==x )
7           {
8               if ( 0==f[n-j] )  { f[n] = 1;  xgen(n+1, n, -x); }
9               f[n] = f[n-j];  xgen(n+1, j, +x);
10          }
11          else
12          {
13              f[n] = f[n-j];  xgen(n+1, j, +x);
14              if ( 0==f[n-j] )  { f[n] = 1;  xgen(n+1, n, -x); }
15          }
16      }
17  }
```

The program [FXT: comb/necklace-gray-demo.cc] computes the binary Lyndon words with the given routine. The ordering obtained has fewer transitions between successive elements but is in general not a Gray code (for up to 6-bit words a Gray code is generated). Figure 17.1-C shows the output with 8-bit Lyndon words. The first $2^{\lfloor n/2 \rfloor} - 1$ Lyndon words of length n are in Gray code order. The number X_n of additional transitions of the length-n Lyndon words is, for $n \le 30$, shown in figure 17.1-D.

17.1.5 An order with at most three changes per transition

An algorithm to generate necklaces in an order such that at most 3 elements change with each update is given in [270]. The recursion can be given as (corrected and shortened) [FXT: comb/necklace-gray3-demo.cc]:

```
1   long *f; // data in f[1..m],  f[0] = 0
2   long N; // word length
3   int k;  // k-ary necklaces, k==sigma in the paper
4
5   void gen3(int z, int t, int j)
6   {
```

```
 1:  .1111111          13:  ...1...1         25:  ..1.1111
 2:  .111.111          14:  ...1.1.1         26:  ..1.11.1
 3:  .11.1111 <<+1     15:  ...1.111         27:  ..1.1.11 <<+1
 4:  .1.1.111 <<+2     16:  .....111         28:  ..1..1.1 <<+2
 5:  .1.1.1.1          17:  .....1.1         29:  ..1..111
 6:  .1.11.11 <<+2     18:  .......1         30:  ..11.111
 7:  .1.11111          19:  ........         31:  ..11.1.1
 8:  ...11111          20:  ......11 <<+1    32:  ..11..11 <<+1
 9:  ...111.1          21:  ....1.11         33:  ..111.11
10:  ...11..1          22:  ....1..1         34:  ..1111.1 <<+1
11:  ...11.11          23:  ....11.1         35:  ..111111
12:  ...1..11          24:  ....1111         36:  11111111 <<+1
```

Figure 17.1-E: The 30 binary 8-bit necklaces in an order with at most 3 changes per transition. Transitions where more than one bit changes are marked with a '<<'.

$n:$	X_n	$n:$	X_n	$n:$	X_n	$n:$	X_n	$n:$	X_n
1:	0	7:	6	13:	200	19:	6462	25:	239008
2:	1	8:	12	14:	360	20:	11722	26:	441370
3:	2	9:	20	15:	628	21:	21234	27:	816604
4:	2	10:	38	16:	1128	22:	38754	28:	1515716
5:	2	11:	64	17:	1998	23:	70770	29:	2818928
6:	4	12:	116	18:	3606	24:	129970	30:	5256628

Figure 17.1-F: Excess (with respect to Gray code) of number of bits changed.

```
7       if ( t > N ) { visit(j); }
8       else
9       {
10          if ( (z&1)==0 ) // z (number of elements ==(k-1)) is even?
11          {
12              for (int i=f[t-j]; i<=k-1; ++i)
13              {
14                  f[t] = i;
15                  gen3( z+(i!=k-1), t+1, (i!=f[t-j]?t:j) );
16              }
17          }
18          else
19          {
20              for (int i=k-1; i>=f[t-j]; --i)
21              {
22                  f[t] = i;
23                  gen3( z+(i!=k-1), t+1, (i!=f[t-j]?t:j) );
24              }
25          }
26      }
27  }
```

The variable `z` counts the number of maximal elements. The output with length-8 binary necklaces is shown in figure 17.1-E. Selecting the necklaces from the reversed list of complemented Gray codes of the n-bit binary words produces the same list.

17.1.6 Binary necklaces of length 2^n via Gray-cycle leaders *

The algorithm for the generation of cycle leaders for the Gray code permutation given section 2.8.1 on page 106 and relation 1.19-10c on page 54, written as

$$S_k\, Y\, x \quad = \quad Y\, g^k\, x \tag{17.1-1}$$

(Y is the yellow code, or bit-wise Reed-Muller transform) allow us to generate the necklaces of length 2^n: The cyclic shifts of $Y\, x$ are equal to $Y\, g^k\, x$ for $k = 0, \ldots, l-1$ where l is the cycle length. Figure 17.1-G shows the correspondence between cycles of the Gray code permutation and cyclic shifts, it was generated with the program [FXT: comb/necklaces-via-gray-leaders-demo.cc].

```
 k = 7: 16 cycles of length= 8
L=  1.......  [ 1.......  ]
L=  1......1  [ .1111111  ]
L=  1.....1.  [ ..1.1.1.  ]
L=  1.....11  [ 11.1.1.1  ]
L=  1....1..  [ .1..11..  ]
L=  1....1.1  [ 1.11..11  ]
L=  1....11.  [ 111..11.  ]
L=  1....111  [ ...11..1  ]
L=  1..1....  [ .111....  ]
L=  1..1...1  [ 1...1111  ]
L=  1..1..1.  [ 11.11.1.  ]
L=  1..1..11  [ ..1..1.1  ]
L=  1..1.1..  [ 1.1111..  ]
L=  1..1.1.1  [ .1....11  ]
L=  1..1.11.  [ ...1.11.  ]
L=  1..1.111  [ 111.1..1  ]

L=  1..1.11.  [ ...1.11.  ]
--> 11.111.1  [ ....1.11  ]
--> 1.11..11  [ 1....1.1  ]
--> 111.1.1.  [ 11....1.  ]
--> 1..11111  [ .11....1  ]
--> 11.1....  [ 1.11....  ]
--> 1.111...  [ .1.11...  ]
--> 111..1..  [ ..1.11..  ]

L=  1..1.111  [ 111.1..1  ]
--> 11.111..  [ 1111.1..  ]
--> 1.11..1.  [ .1111.1.  ]
--> 111.1.11  [ ..1111.1  ]
--> 1..1111.  [ 1..1111.  ]
--> 11.1...1  [ .1..1111  ]
--> 1.111..1  [ 1.1..111  ]
--> 111..1.1  [ 11.1..11  ]
```

Figure 17.1-G: Left: the cycle leaders (minima) L of the Gray code permutation where is highest bit has index 7 and their bit-wise Reed-Muller transforms $Y(L)$. Right: the last two cycles and the transforms of their elements.

If no better algorithm for the cyclic leaders of the Gray code permutation was known we could generate them as $Y^{-1}(N) = Y(N)$ where N are the necklaces of length 2^n. The same idea, and relation 1.19-11b on page 54, give the relation

$$S_k\, B\, x \quad = \quad B\, e^{-k}\, x \tag{17.1-2}$$

where B is the blue code and e the reversed Gray code.

17.1.7 Binary necklaces via cyclic shifts and complements *

```
  n = 3
1:  ..1
2:  .11
3:  111

  n = 4
1:  ...1
2:  ..11
3:  .111
4:  1111
5:  .1.1

  n = 5
1:  ....1
2:  ...11
3:  ..111
4:  .1111
5:  11111
6:  ..1.1
7:  .1.11

  n = 6
1:  .....1
2:  ....11
3:  ...111
4:  ..1111
5:  .11111
6:  111111
7:  ..11.1
8:  .11.11
9:  ...1.1
10: ..1.11
11: .1.111
12: .1.1.1
13: ..1..1

  n = 7
1:  ......1
2:  .....11
3:  ....111
4:  ...1111
5:  ..11111
6:  .111111
7:  1111111
8:  ..111.1
9:  ...11.1
10: ..11.11
11: .11.111
12: ....1.1
13: ...1.11
14: ..1.111
15: .1.1111
16: ..1.1.1
17: .1.1.11
18: ...1..1
19: ..1..11

  n = 8
1:  .......1
2:  ......11
3:  .....111
4:  ....1111
5:  ...11111
6:  ..111111
7:  .1111111
8:  11111111
9:  ..1111.1
10: ...111.1
11: ..111.11
12: .111.111
13: ....11.1
14: ...11.11
15: ..11.111
16: .11.1111
17: ..11.1.1
18: ...11..1

  [n=8 cont.]
19:  ..11..11
20:  .....1.1
21:  ....1.11
22:  ...1.111
23:  ..1.1111
24:  .1.11111
25:  ..1.11.1
26:  .1.11.11
27:  ...1.1.1
28:  ..1.1.11
29:  .1.1.111
30:  .1.1.1.1
31:  ....1..1
32:  ...1..11
33:  ..1..111
34:  ..1..1.1
35:  ...1...1
```

Figure 17.1-H: Nonzero binary necklaces of lengths $n = 3, 4, \ldots, 8$ as generated by the shift and complement algorithm.

A recursive algorithm to generate all nonzero binary necklaces via cyclic shifts and complements of the lowest bit is described in [223]. An implementation of the method is given in [FXT: comb/necklace-sigma-tau-demo.cc]:

```
1  inline ulong sigma(ulong x)  { return bit_rotate_left(x, 1, n); }
2  inline ulong tau(ulong x)  { return  x ^ 1; }
3
4  void search(ulong y)
5  {
6      visit(y);
7      ulong t = y;
```

```
 8       while ( 1 )
 9       {
10           t = sigma(t);
11           ulong x = tau(t);
12           if ( (x&1) && (x == bit_cyclic_min(x, n)) )  search(x);
13           else  break;
14       }
15   }
```

The initial call is `search(1)`. The generated ordering for lengths $n = 3, 4, \ldots, 8$ is shown in figure 17.1-H.

17.2 The number of binary necklaces

n :	N_n	n :	N_n	n :	N_n	n :	N_n
1:	2	11:	188	21:	99880	31:	69273668
2:	3	12:	352	22:	190746	32:	134219796
3:	4	13:	632	23:	364724	33:	260301176
4:	6	14:	1182	24:	699252	34:	505294128
5:	8	15:	2192	25:	1342184	35:	981706832
6:	14	16:	4116	26:	2581428	36:	1908881900
7:	20	17:	7712	27:	4971068	37:	3714566312
8:	36	18:	14602	28:	9587580	38:	7233642930
9:	60	19:	27596	29:	18512792	39:	14096303344
10:	108	20:	52488	30:	35792568	40:	27487816992

Figure 17.2-A: The number of binary necklaces for $n \leq 40$.

n :	L_n	n :	L_n	n :	L_n	n :	L_n
1:	2	11:	186	21:	99858	31:	69273666
2:	1	12:	335	22:	190557	32:	134215680
3:	2	13:	630	23:	364722	33:	260300986
4:	3	14:	1161	24:	698870	34:	505286415
5:	6	15:	2182	25:	1342176	35:	981706806
6:	9	16:	4080	26:	2580795	36:	1908866960
7:	18	17:	7710	27:	4971008	37:	3714566310
8:	30	18:	14532	28:	9586395	38:	7233615333
9:	56	19:	27594	29:	18512790	39:	14096302710
10:	99	20:	52377	30:	35790267	40:	27487764474

Figure 17.2-B: The number of binary Lyndon words for $n \leq 40$.

The number of binary necklaces of length n equals

$$N_n \;=\; \frac{1}{n}\sum_{d\backslash n}\varphi(d)\,2^{n/d} \;=\; \frac{1}{n}\sum_{j=1}^{n}2^{\gcd(j,n)} \tag{17.2-1}$$

Replace 2 by k to get the number for k different sorts of beads (possible values at each digit). The values for $n \leq 40$ are shown in figure 17.2-A. The sequence is entry A000031 of [240].

The number of Lyndon words (aperiodic necklaces) equals

$$L_n \;=\; \frac{1}{n}\sum_{d\backslash n}\mu(d)\,2^{n/d} \;=\; \frac{1}{n}\sum_{d\backslash n}\mu(n/d)\,2^{d} \tag{17.2-2}$$

The Möbius function μ is defined in relation 35.1-6 on page 679. The values for $n \leq 40$ are given in figure 17.2-B. The sequence is entry A001037 of [240]. For prime $n = p$ we have $L_p = N_p - 2$ and

$$L_p \;=\; \frac{2^p - 2}{p} \;=\; \frac{1}{p} \sum_{k=1}^{p-1} \binom{p}{k} \tag{17.2-3}$$

The latter form transpires that there are exactly $\binom{p}{k}/p$ Lyndon words with k ones for $1 \leq k \leq p-1$. The difference of two is due to the necklaces that consist of all zeros or ones. The number of irreducible binary polynomials (see section 38.6 on page 837) of degree n also equals L_n. For the equivalence between necklaces and irreducible polynomials see section 38.10 on page 849.

Let d be a divisor of n. There are 2^n binary words of length n, each having some period d that divides n. There are d different shifts of the corresponding word, thereby

$$2^n \;=\; \sum_{d \backslash n} d\, L_d \tag{17.2-4}$$

Möbius inversion gives relation 17.2-2. The necklaces of length n and period d are obtained by concatenation of n/d Lyndon words of length d, so

$$N_n \;=\; \sum_{d \backslash n} L_d \tag{17.2-5}$$

We note the relations (see section 35.1.4 on page 681)

$$(1 - 2\,y) \;=\; \prod_{k=1}^{\infty} (1 - y^k)^{L_k} \tag{17.2-6a}$$

$$\sum_{k=1}^{\infty} L_k\, y^k \;=\; \sum_{k=1}^{\infty} \frac{-\,\mu(k)}{k} \log\left(1 - 2\,y^k\right) \tag{17.2-6b}$$

Defining

$$\eta_B(x) \;:=\; \prod_{k=1}^{\infty} \left(1 - B\,y^k\right) \tag{17.2-7a}$$

we have

$$\eta_2(x) \;=\; \prod_{k=1}^{\infty} (1 - y^k)^{N_k} \tag{17.2-7b}$$

$$\eta_2(x) \;=\; \prod_{k=1}^{\infty} \eta_1(y^k)^{L_k} \tag{17.2-7c}$$

17.3 The number of binary necklaces with fixed content

Let $N_{\binom{n}{n_0}}$ be the number of binary length-n necklaces with exactly n_0 zeros (and $n_1 = n - n_0$ ones). We call these *necklaces with fixed density*. One has

$$N_{\binom{n}{n_0}} \;=\; \frac{1}{n} \sum_{j \backslash g} \varphi(j) \binom{n/j}{d/j} \tag{17.3-1}$$

where $g = \gcd(n, n_0)$. One has the symmetry relation $N_{\binom{n}{n_0}} = N_{\binom{n}{n-n_0}} = N_{\binom{n}{n_1}}$. A table of small values is given in figure 17.3-A.

n:	N_n	$N_{\binom{n}{0}}$	$N_{\binom{n}{1}}$	$N_{\binom{n}{2}}$	$N_{\binom{n}{3}}$	$N_{\binom{n}{4}}$	$N_{\binom{n}{5}}$	$N_{\binom{n}{6}}$	$N_{\binom{n}{7}}$	$N_{\binom{n}{8}}$	$N_{\binom{n}{9}}$	$N_{\binom{n}{10}}$
1:	2	1	1									
2:	3	1	1	1								
3:	4	1	1	1	1							
4:	6	1	1	2	1	1						
5:	8	1	1	2	2	1	1					
6:	14	1	1	3	4	3	1	1				
7:	20	1	1	3	5	5	3	1	1			
8:	36	1	1	4	7	10	7	4	1	1		
9:	60	1	1	4	10	14	14	10	4	1	1	
10:	108	1	1	5	12	22	26	22	12	5	1	1
11:	188	1	1	5	15	30	42	42	30	15	5	1
12:	352	1	1	6	19	43	66	80	66	43	19	6
13:	632	1	1	6	22	55	99	132	132	99	55	22
14:	1182	1	1	7	26	73	143	217	246	217	143	73
15:	2192	1	1	7	31	91	201	335	429	429	335	201
16:	4116	1	1	8	35	116	273	504	715	810	715	504
17:	7712	1	1	8	40	140	364	728	1144	1430	1430	1144
18:	14602	1	1	9	46	172	476	1038	1768	2438	2704	2438
19:	27596	1	1	9	51	204	612	1428	2652	3978	4862	4862
20:	52488	1	1	10	57	245	776	1944	3876	6310	8398	9252

Figure 17.3-A: The number of binary necklaces with a prescribed number of ones.

n:	L_n	$L_{\binom{n}{0}}$	$L_{\binom{n}{1}}$	$L_{\binom{n}{2}}$	$L_{\binom{n}{3}}$	$L_{\binom{n}{4}}$	$L_{\binom{n}{5}}$	$L_{\binom{n}{6}}$	$L_{\binom{n}{7}}$	$L_{\binom{n}{8}}$	$L_{\binom{n}{9}}$	$L_{\binom{n}{10}}$
1:	2	1	1									
2:	1	0	1	0								
3:	2	0	1	1	0							
4:	3	0	1	1	1	0						
5:	6	0	1	2	2	1	0					
6:	9	0	1	2	3	2	1	0				
7:	18	0	1	3	5	5	3	1	0			
8:	30	0	1	3	7	8	7	3	1	0		
9:	56	0	1	4	9	14	14	9	4	1	0	
10:	99	0	1	4	12	20	25	20	12	4	1	0
11:	186	0	1	5	15	30	42	42	30	15	5	1
12:	335	0	1	5	18	40	66	75	66	40	18	5
13:	630	0	1	6	22	55	99	132	132	99	55	22
14:	1161	0	1	6	26	70	143	212	245	212	143	70
15:	2182	0	1	7	30	91	200	333	429	429	333	200
16:	4080	0	1	7	35	112	273	497	715	800	715	497
17:	7710	0	1	8	40	140	364	728	1144	1430	1430	1144
18:	14532	0	1	8	45	168	476	1026	1768	2424	2700	2424
19:	27594	0	1	9	51	204	612	1428	2652	3978	4862	4862
20:	52377	0	1	9	57	240	775	1932	3876	6288	8398	9225

Figure 17.3-B: The number of binary Lyndon words with a prescribed number of ones.

Let $L_{\binom{n}{n_0}}$ be the number of binary length-n Lyndon words with exactly n_0 ones (*Lyndon words with fixed density*), then

$$L_{\binom{n}{n_0}} = \frac{1}{n} \sum_{j \backslash g} \mu(j) \binom{n/j}{d/j} \tag{17.3-2}$$

where $g = \gcd(n, n_0)$. The symmetry relation is the same as for $N_{\binom{n}{n_0}}$. A table of small values is given in figure 17.3-B.

Let $N(n_0, n_1, \ldots, n_{k-1})$ be the number of k-symbol length-n necklaces with n_j occurrences of symbol j. For the number of such *necklaces with fixed content* we have ($n = \sum_{j<s} n_j$ and):

$$N(n_0, n_1, \ldots, n_{k-1}) = \frac{1}{n} \sum_{d \backslash g} \varphi(d) \frac{(n/d)!}{(n_0/d)! \cdots (n_{k-1}/d)!} \tag{17.3-3}$$

where $g = \gcd(n_0, n_1, \ldots, n_{k-1})$. The equivalent formula for the *Lyndon words with fixed content* is

$$L(n_0, n_1, \ldots, n_{k-1}) = \frac{1}{n} \sum_{d \backslash g} \mu(d) \frac{(n/d)!}{(n_0/d)! \cdots (n_{k-1}/d)!} \tag{17.3-4}$$

where $g = \gcd(n_0, n_1, \ldots, n_{k-1})$. The relations were taken from [224] and [232] which also give efficient algorithms for the generation of necklaces and Lyndon words with fixed density and content, respectively. The number of strings with fixed content is a *multinomial coefficient*, see relation 11.2-1a on page 290. A method for the generation of all necklaces with forbidden substrings is given in [225].

Chapter 18

Hadamard and conference matrices

The $2^k \times 2^k$-matrices corresponding to the Walsh transforms (see chapter 22 on page 447) are special cases of so-called Hadamard matrices. Such matrices also exist for certain sizes $n \times n$ for n not a power of two. We give construction schemes for Hadamard matrices that come from the theory of finite fields.

If we denote the transform matrix for an N-point Walsh transform by $\mathbf{H}$, then

$$\mathbf{H}\,\mathbf{H}^T \quad = \quad N \text{ id} \tag{18.0-1}$$

where id is the unit matrix. The matrix $\mathbf{H}$ is orthogonal (up to normalization) and its determinant equals

$$\det(\mathbf{H}) \quad = \quad \det\left(\mathbf{H}\,\mathbf{H}^T\right)^{1/2} \; = \; N^{N/2} \tag{18.0-2}$$

Further, all entries are either $+1$ or -1. An orthogonal matrix with these properties is called a *Hadamard matrix*. We know that for $N = 2^n$ we always can find such a matrix. For $N = 2$ we have

$$\mathbf{H}_2 \quad = \quad \begin{bmatrix} +1 & +1 \\ +1 & -1 \end{bmatrix} \tag{18.0-3}$$

and we can use the Kronecker product (see section 22.3 on page 451) to construct $\mathbf{H}_{2N}$ from $\mathbf{H}_N$ via

$$\mathbf{H}_n \quad = \quad \begin{bmatrix} +\mathbf{H}_{N/2} & +\mathbf{H}_{N/2} \\ +\mathbf{H}_{N/2} & -\mathbf{H}_{N/2} \end{bmatrix} \; = \; \mathbf{H}_2 \otimes \mathbf{H}_{N/2} \tag{18.0-4}$$

The problem of determining Hadamard matrices (especially for N not a power of two) comes from combinatorics. Hadamard matrices of size $N \times N$ can only exist if N equals 1, 2, or $4\,k$.

18.1 Hadamard matrices via LFSR

We start with a construction for certain Hadamard matrices for N a power of two that uses m-sequences that are created by shift registers (see section 39.1 on page 857 and section 39.4 on page 862). Figure 18.1-A shows three Hadamard matrices that were constructed as follows:

1. Choose $N = 2^n$ and create a maximum length binary shift register sequence S of length $N - 1$.
2. Make S signed, that is, replace all ones by -1 and all zeros by $+1$.
3. The $N \times N$ matrix $\mathbf{H}$ is obtained by filling the first row and the first column with ones and filling the remaining entries with cyclical copies of s: for $r = 1, 2, \ldots N-1$ and $c = 1, 2, \ldots N-1$ set $\mathbf{H}_{r,c} := S_{c-r+1 \bmod N-1}$.

```
Signed SRS:
   - + + + - + + - - + - + - - -
Hadamard matrix H:
 + + + + + + + + + + + + + + + +
 + - + + + - + + - - + - + - - -
 + - - + + + - + + - - + - + - -
 + - - - + + + - + + - - + - + -
 + - - - - + + + - + + - - + - +
 + + - - - - + + + - + + - - + -
 + - + - - - - + + + - + + - - +
 + + - + - - - - + + + - + + - -
 + - + - + - - - - + + + - + + -
 + - - + - + - - - - + + + - + +
 + + - - + - + - - - - + + + - +
 + + + - - + - + - - - - + + + -
 + - + + - - + - + - - - - + + +
 + + - + + - - + - + - - - - + +
 + + + - + + - - + - + - - - - +
 + + + + - + + - - + - + - - - -

Signed SRS:
   - + + - + - -
Hadamard matrix H:
 + + + + + + + +
 + - + + - + - -
 + - - + + - + -
 + - - - + + - +
 + + - - - + + -
 + - + - - - + +
 + + - + - - - +
 + + + - + - - -

Signed SRS:
   - + -
Hadamard matrix H:
 + + + +
 + - + -
 + - - +
 + + - -
```

Figure 18.1-A: Hadamard matrices created with binary shift register sequences (SRS) of maximum length. Only the sign of the entries is given, all entries are ± 1.

For given n we can obtain as many different Hadamard matrices as there are m-sequences.

The matrices in figure 18.1-A where produced with the program [FXT: comb/hadamard-srs-demo.cc].

```
#include "bpol/lfsr.h" // class lfsr
#include "aux1/copy.h"  // copy_cyclic()

#include "matrix/matrix.h"  // class matrix
typedef matrix<int>  Smat;  // matrix with integer entries

  [--snip--]
    ulong n = 5;
    ulong N = 1UL << n;
  [--snip--]

    // --- create signed SRS:
    int vec[N-1];
    lfsr S(n);
    for (ulong k=0; k<N-1; ++k)
    {
        ulong x = 1UL & S.get_a();
        vec[k] = ( x ? -1 : +1 );
        S.next();
    }

    // --- create Hadamard matrix:
    Smat H(N,N);
    for (c=0; c<N; ++c)  H.set(0, c, +1);  // first row = [1,1,1,...,1]
    for (ulong r=1; r<N; ++r)
    {
        H.set(r, 0, +1);  // first column = [1,1,1,...,1]^T
        copy_cyclic(vec, H.rowp_[r]+1, N-1, N-r);
    }
  [--snip--]
```

The function `copy_cyclic()` is defined in [FXT: aux1/copy.h]:

```
template <typename Type>
inline void copy_cyclic(const Type *src, Type *dst, ulong n, ulong s)
// Copy array src[] to dst[]
// starting from position s in src[]
// wrap around end of src[]  (src[n-1])
//
// src[] is assumed to be of length n
// dst[] must be length n at least
//
// Equivalent to:  { copy(src, dst, n); rotate_right(dst, n, s)}
{
    ulong k = 0;
    while ( s<n )  dst[k++] = src[s++];

    s = 0;
```

```
16      while ( k<n )  dst[k++] = src[s++];
17  }
```

If we define the matrix $\mathbf{X}$ to be the $(N-1) \times (N-1)$ block of $\mathbf{H}$ obtained by deleting the first row and column then

$$\mathbf{X}\,\mathbf{X}^T \quad = \quad \begin{bmatrix} N-1 & -1 & -1 & \cdots & -1 \\ -1 & N-1 & -1 & \cdots & -1 \\ -1 & -1 & N-1 & \cdots & -1 \\ \vdots & \vdots & \cdots & \ddots & \vdots \\ -1 & -1 & -1 & \cdots & N-1 \end{bmatrix} \tag{18.1-1}$$

Equivalently, for the (cyclic) auto correlation of S (see section 39.5 on page 863):

$$\sum_{k=0}^{L-1} S_k\, S_{k+\tau \bmod L} \quad = \quad \begin{cases} +L & \text{if} \quad \tau = 0 \\ -1 & \text{else} \end{cases} \tag{18.1-2}$$

where $L = N - 1$ is the length of the sequence.

A alternative way to obtain Hadamard matrices of dimension 2^n is to use the signs in the multiplication table for hypercomplex numbers described in section 37.16 on page 810.

18.2 Hadamard matrices via conference matrices

```
Quadratic characters modulo 13:
   0 + - + + - - - - + + - +
 14x14 conference matrix C:
  0 + + + + + + + + + + + +
  + 0 + - + + - - - - + + - +
  + + 0 + - + + - - - - + + -
  + - + 0 + - + + - - - - + +
  + + - + 0 + - + + - - - - +
  + + + - + 0 + - + + - - -
  + - + + - + 0 + - + + - -
  + - - + + - + 0 + - + + -
  + - - - + + - + 0 + - + +
  + - - - - + + - + 0 + - +
  + + - - - - + + - + 0 + -
  + + + - - - - + + - + 0 +
  + - + + - - - - + + - + 0
  + + - + + - - - - + + - +
```

```
Quadratic characters modulo 11:
   0 + - + + + - - - + -
12x12 conference matrix C:
  0 + + + + + + + + + + +
  - 0 + - + + + - - - + -
  - - 0 + - + + + - - - +
  - + - 0 + - + + + - - -
  - - + - 0 + - + + + - -
  - - - + - 0 + - + + + -
  - - - - + - 0 + - + + +
  - + - - - + - 0 + - + +
  - + + - - - + - 0 + - +
  - + + + - - - + - 0 + -
  - - + + + - - - + - 0 +
  - + - + + + - - - + - 0
```

Figure 18.2-A: Two Conference matrices, the entries not on the diagonal are ± 1 and only the sign is given. The left is a symmetric 14×14 matrix ($13 \equiv 1 \bmod 4$), the right is a antisymmetric 12×12 matrix ($11 \equiv 3 \bmod 4$). Replacing all diagonal elements of the right matrix with $+1$ gives a 12×12 Hadamard matrix.

A *conference matrix* $\mathbf{C}_Q$ is a $Q \times Q$ matrix with zero diagonal and all other entries ± 1 so that

$$\mathbf{C}\,\mathbf{C}^T \quad = \quad (Q-1)\ \mathrm{id} \tag{18.2-1}$$

An algorithm for the construction of $\mathbf{C}_q$ for $Q = q + 1$ where q is an odd prime:

1. Create a length-q array S with entries $S_k \in \{-1,\, 0,\, +1\}$ as follows: set $S_0 = 0$ and, for $1 \le k < q$ set $S_k = +1$ if k is a square modulo q, $S_k = -1$ else.
2. Set $y = 1$ if $q \equiv 1 \bmod 4$, else $y = -1$ (then $q \equiv 3 \bmod 4$).
3. Create a $Q \times Q$ matrix $\mathbf{C}$ as follows: set $\mathbf{C}_{0,0} = 0$ and $\mathbf{C}_{0,k} = +1$ for $1 \le k < Q$ (first row). Set $\mathbf{C}_{k,0} = y$ for $1 \le k < Q$ (first column). Fill the remaining entries with cyclical copies of S: for $1 \le r < q$ and $1 \le c < q$ set $\mathbf{C}_{r,c} = S_{c-r+1 \bmod q}$.

The quantity y tells us whether $\mathbf{C}$ is symmetric ($y = +1$) or antisymmetric ($y = -1$).

If $\mathbf{C}_Q$ is antisymmetric then $H_Q = \mathbf{C}_Q + \mathrm{id}$ is a Hadamard matrix. For example, replacing all zeros in the 12×12 matrix in figure 18.2-A by $+1$ gives a 12×12 Hadamard matrix..

```
12x12 Hadamard matrix H:
 + + + + + + - + + + + +
 + + + - - + + - + - - +
 + + + + - - + + - + - -
 + - + + + - + - + - + -
 + - - + + + + - - + - +
 + + - - + + + + - - + -
 - + + + + + - - - - - -
 + - + - - + - - - + + -
 + + - + - - - - - - + +
 + - + - + - - + - - - +
 + - - + - + - + + - - -
 + + - - + - - - + + - -

Quadratic characters modulo 5:
   0 + - - +
6x6 conference matrix C:
 0 + + + + +
 + 0 + - - +
 + + 0 + - -
 + - + 0 + -
 + - - + 0 +
 + + - - + 0
```

Figure 18.2-B: A 12×12 Hadamard matrix (left) created from the symmetric 6×6 conference matrix on the right.

If $\mathbf{C}_Q$ is symmetric then a $2Q \times 2Q$ Hadamard matrix is given by

$$\mathbf{H}_{2Q} \quad := \quad \begin{bmatrix} +\mathrm{id} + \mathbf{C} & -\mathrm{id} + \mathbf{C} \\ -\mathrm{id} + \mathbf{C} & -\mathrm{id} - \mathbf{C} \end{bmatrix} \tag{18.2-2}$$

Figure 18.2-B shows a 12×12 Hadamard matrix that was created using this formula.

The program [FXT: comb/conference-quadres-demo.cc] outputs for a given q the $Q \times Q$ conference matrix and the corresponding Hadamard matrix:

```
#include "mod/numtheory.h" // kronecker()
#include "matrix/matrix.h"  // class matrix
#include "aux1/copy.h"  // copy_cyclic()

  [--snip--]
    int y = ( 1==q%4 ? +1 : -1 );
    ulong Q = q+1;
  [--snip--]
    // --- create table of quadratic characters modulo q:
    int vec[q];  fill<int>(vec, q, -1);  vec[0] = 0;
    for (ulong k=1; k<(q+1)/2; ++k)  vec[(k*k)%q] = +1;
  [--snip--]
    // --- create Q x Q conference matrix:
    Smat C(Q,Q);
    C.set(0,0, 0);
    for (ulong c=1; c<Q; ++c)  C.set(0, c, +1);  // first row = [1,1,1,...,1]
    for (ulong r=1; r<Q; ++r)
    {
        C.set(r, 0, y);  // first column = +-[1,1,1,...,1]^T
        copy_cyclic(vec, C.rowp_[r]+1, q, Q-r);
    }
  [--snip--]
    // --- create a N x N Hadamard matrix:
    ulong N = ( y<0 ? Q : 2*Q );
    Smat H(N,N);
    if ( N==Q )
    {
        copy(C, H);
        H.diag_add_val(1);
    }
    else
    {
        Smat K2(2,2);  K2.fill(+1); K2.set(1,1, -1);  // K2 = [+1,+1; +1,-1]
        H.kronecker(K2, C);  // Kronecker product of matrices
        for (ulong k=0; k<Q; ++k)  // adjust diagonal of submatrices
        {
            ulong r, c;
            r=k;   c=k;   H.set(r,c, H.get(r,c)+1);
            r=k;   c=k+Q; H.set(r,c, H.get(r,c)-1);
            r=k+Q; c=k;   H.set(r,c, H.get(r,c)-1);
```

```
41              r=k+Q; c=k+Q; H.set(r,c, H.get(r,c)-1);
42          }
43      }
44    [--snip--]
```

If both $\mathbf{H}_a$ and $\mathbf{H}_b$ are Hadamard matrices (of dimensions a and b, respectively) then their Kronecker product $\mathbf{H}_{ab} = \mathbf{H}_a \otimes \mathbf{H}_b$ is again a Hadamard matrix:

$$\mathbf{H}_{ab}\,\mathbf{H}_{ab}^T \;=\; (\mathbf{H}_a \otimes \mathbf{H}_b)\,(\mathbf{H}_a \otimes \mathbf{H}_b)^T \;=^*\; (\mathbf{H}_a \otimes \mathbf{H}_b)\,\left(\mathbf{H}_a^T \otimes \mathbf{H}_b^T\right) \;= \tag{18.2-3a}$$

(the starred equality uses relation 22.3-11a on page 452)

$$= \;\left(\mathbf{H}_a\,\mathbf{H}_a^T\right) \otimes \left(\mathbf{H}_b \otimes \mathbf{H}_b^T\right) \;=^*\; (a\ \mathrm{id}) \otimes (b\ \mathrm{id}) \;=\; a\,b\ \mathrm{id} \tag{18.2-3b}$$

(the starred equality uses relation 22.3-10a on page 452).

18.3 Conference matrices via finite fields

The algorithm for odd primes q can be modified to work also for powers of odd primes. Then one has to work with the finite fields $\mathrm{GF}(q^n)$. The entries $\mathbf{C}_{r+1,c+1}$ for $r = 0, 1, \ldots, q^n - 1$ and $c = 0, 1, \ldots, q^n - 1$ have to be the quadratic character of $z_r - z_c$ where $z_0, z_1, \ldots, z_{q^n-1}$ are the elements in $\mathrm{GF}(q^n)$ in some (fixed) order.

We implement the algorithm in pari/gp. Firstly, we give two simple routines that map the elements $z_i \in \mathrm{GF}(q^n)$ (represented as polynomials modulo q) to the numbers $0, 1, \ldots, q^n - 1$. The polynomial $p(x) = c_0 + c_1\,x + \ldots + c_{n-1}\,x^{n-1}$ is mapped to $N = c_0 + c_1\,q + \ldots + c_{n-1}\,q^{n-1}$.

```
1   pol2num(p,q)=
2   \\ Return number for polynomial p.
3   {
4       local(t, n);  n=0;
5       for (k=0, poldegree(p),
6           t = polcoeff(p, k);
7           t *= Mod(1,q);
8           t = component(t, 2);
9           t *= q^k;
10          n += t;
11      );
12      return( n );
13  }
```

The inverse routine is

```
1   num2pol(n,q)=
2   \\ Return polynomial for number n.
3   {
4       local(p, mq, k);
5       p = Pol(0,'x);
6       k = 0;
7       while ( 0!=n,
8           mq = n % q;
9           p += mq * ('x)^k;
10          n -= mq;
11          n \= q;
12          k++;
13      );
14      return( p );
15  }
```

The quadratic character of an element z can be determined by computing $z^{(q^n-1)/2}$ modulo the field polynomial. The result will be zero for $z = 0$, else ± 1. The following routine determines the character of the difference of two elements as required for the computation of conference matrices:

```
1   getquadchar_n(n1, n2, q, fp, n)=
2   \\ Return the quadratic character of (n2-n1) in GF(q^n)
3   \\ Powering method
4   {
```

```
5       local(p1, p2, d, nd, sc);
6       if ( n1==n2, return(0) );
7       p1 = num2pol(n1, q);
8       p2 = num2pol(n2, q);
9       d = Mod(1,q)* (p2-p1);
10      d = Mod(d,fp)^((q^n-1)/2);
11      d = component(d, 2);
12      if ( Mod(1,q)==d,  sc=+1, sc=-1 );
13      return( sc );
14  }
```

The input are two numbers that are mapped to the corresponding field elements.

In order to reduce the computational work we create a table of the quadratic characters for later lookup:

```
1   quadcharvec(fp, q)=
2   \\ Return a table of quadratic characters in GF(q^n)
3   \\ fp is the field polynomial.
4   {
5       local(n, qn, sv, pl);
6       n=poldegree(fp);
7       qn=q^n-1;
8       sv=vector(qn+1, j, -1);
9       sv[1] = 0;
10      for (k=1, qn,
11          pl = num2pol(k,q);
12          pl = Mod(Mod(1,q)*pl, fp);
13          sq = pl * pl;
14          sq = component(sq, 2);
15          i = pol2num( sq, q );
16          sv[i+1] = +1;
17      );
18      return( sv );
19  }
```

Using this table we can compute the quadratic characters of the difference of two elements in a more efficient manner:

```
1   getquadchar_v(n1, n2, q, fp, sv)=
2   \\ Return the quadratic character of (n2-n1) in GF(q^n)
3   \\ Table lookup method
4   {
5       local(p1, p2, d, nd, sc);
6       if ( n1==n2, return(0) );
7       p1 = num2pol(n1, q);
8       p2 = num2pol(n2, q);
9       d = (p2-p1) % fp;
10      nd = pol2num(d, q);
11      sc = sv[nd+1];
12      return( sc );
13  }
```

Now we can compute conference matrices:

```
1   matconference(q, fp, sv)=
2   \\ Return a QxQ conference matrix.
3   \\ q  an odd prime.
4   \\ fp an irreducible polynomial modulo q.
5   \\ sv table of quadratic characters in GF(q^n)
6   \\   where n is the degree of fp.
7   {
8       local(y, Q, C, n);
9       n = poldegree(fp);
10      Q=q^n+1;
11      C = matrix(Q,Q);
12      for (k=2, Q, C[1,k]=+1);  \\ first row
13      for (k=2, Q, C[k,1]=+1);  \\ first column
14      for (r=2, Q,
15          for (c=2, Q,
16              sc = getquadchar_n(r-2, c-2, q, fp, n);
17  \\            sc = getquadchar_v(r-2, c-2, q, fp, sv);  \\ same result
18              C[r,c] = sc;
19          );
20      );
```

```
21      return( C );
22  }
```

```
q = 3   fp = x^2 + 1   GF(3^2)
Table of quadratic characters:
   0 + + + - - + - -
10x10 conference matrix C:
 0 + + + + + + + + +
 - 0 + + + - - + - -
 - + 0 + - + - - + -
 - + + 0 - - + - - +
 - + - - 0 + + + - -
 - - + - + 0 + - + -
 - - - + + + 0 - - +
 - + - - + - - 0 + +
 - - + - - + - + 0 +
 - - - + - - + + + 0
```

Figure 18.3-A: A 10×10 conference matrix for $q = 3$ and the field polynomial $f = x^2 + 1$.

```
q = 3   fp = x^3 - x + 1   GF(3^3)
Table of quadratic characters:
   0 + - - - - + + + + - + + + - + + + - - - + - + + - + -
28x28 conference matrix C:
 0 + + + + + + + + + + + + + + + + + + + + + + + + + + + +
 - 0 + - - - - + + + + - + + + - - - + - + + - - + - - + -
 - - 0 + - - - + + + + + - - + + - + + + - - - - + - - + +
 - + - 0 - - - + + + - + + + - + + - + - + - + - - + - - -
 - + + + 0 + - - - - + + + + - - + + - - - + - + - - - + -
 - + + + - 0 + - - - - + + + + - - + + - - - + - + - - - +
 - + + + + - 0 - - - + - + - + + + - + + - - - + - + - - -
 - - - - + + + 0 + - + + - + + - + - + - + - - - + - - - +
 - - - - + + + - 0 + - + + + + - - - + - - + + - - + - -
 - - - + + + + + - 0 + - + + - + - + + + - - + - - - + - -
 - - - + - + - - + - 0 + - - - - + + + + - + + + - + + - +
 - + - - - - + - - + 0 + - - - + + + + - - + + + - - + + +
 - - + - + - - + - - + 0 - - - + + + + - + + + - + + - + +
 - - + - - + - + - + + + 0 + - - - - + + - + - + + + - + -
 - - - + - - + + - - - - + + + - 0 + - + + - + + + + - + +
 - + - - + - - - + - - - + + + + - 0 + - + + - + - + + + +
 - + - + + + - + + - - - + - + - - + - 0 + - - - + + + + +
 - + + - - + + - + + + - - - - + - - + - 0 + - - - + + + +
 - - + + + - - + + - - + + - - - + - - - 0 + - - - + + + -
 - + + - + + + - + + - - - + - + - - + + + - 0 + - - - + -
 - + + - + + + - + + - - - + - - - + - - - + + + 0 + - - -
 - - + + + - - + + - - + + - - - - + + + + - 0 + - - - + +
 - + - + - + + + - + + - - - + - + - - + + + - 0 - - - + +
 - + + - + + + - + + - - - + - - - + - - - + + + 0 + - - +
 - - + + + - - + + - - + + - - - - + + + + - 0 + - - - + -
 - + - + - + + + - + + - - - + - + - - + + + - 0 - - - + -
```

If a conference matrix of size c exists then we can create Hadamard matrices of sizes $N = c$ whenever $q^n \equiv 3 \bmod 4$, and $N = 2\,c$ whenever $q^n \equiv 1 \bmod 4$. Further, if Hadamard matrices of sizes N and M exist then a $(N \cdot M) \times (N \cdot M)$ Hadamard matrix can be obtained via the Kronecker product.

The values of $N = 4\,k \leq 2000$ such that this construction does *not* give a $N \times N$ Hadamard matrix are:

```
92, 116, 156, 172, 184, 188, 232, 236, 260, 268, 292, 324, 356, 372,
376, 404, 412, 428, 436, 452, 472, 476, 508, 520, 532, 536, 584,
596, 604, 612, 652, 668, 712, 716, 732, 756, 764, 772, 808, 836,
852, 856, 872, 876, 892, 904, 932, 940, 944, 952, 956, 964, 980,
988, 996, 1004, 1012, 1016, 1028, 1036, 1068, 1072, 1076, 1100,
1108, 1132, 1148, 1168, 1180, 1192, 1196, 1208, 1212, 1220, 1244,
1268, 1276, 1300, 1316, 1336, 1340, 1364, 1372, 1380, 1388, 1396,
1412, 1432, 1436, 1444, 1464, 1476, 1492, 1508, 1528, 1556, 1564,
1588, 1604, 1612, 1616, 1636, 1652, 1672, 1676, 1692, 1704, 1712,
1732, 1740, 1744, 1752, 1772, 1780, 1796, 1804, 1808, 1820, 1828,
1836, 1844, 1852, 1864, 1888, 1892, 1900, 1912, 1916, 1928, 1940,
1948, 1960, 1964, 1972, 1976, 1992
```

This is sequence A046116 of [240]. It can be obtained by starting with a list of all numbers of the form $4\,k$ and deleting all values $k = 2^a\,(q + 1)$ where q is a power of an odd prime. Constructions for Hadamard matrices for numbers of certain forms are known, see [186]. Whether Hadamard matrices exist for all values $N = 4\,k$ is an open problem. A readable source about constructions for Hadamard matrices is [243].

Chapter 19

Searching paths in directed graphs

We describe how certain combinatorial structures can be represented as paths or cycles in a directed graph. As an example consider Gray codes of n-bit binary words: we are looking for sequences of all 2^n binary words such that only one bit changes between two successive words. A convenient representation of the search space is that of a graph. The nodes are the binary words and an edge is drawn between two nodes if the node's values differ by exactly one bit. Every path that visits all nodes of that graph corresponds to a Gray code. If the path is a cycle then a Gray cycle was found.

In general we can, depending on the size of the problem,

1. try to find at least one object
2. generate all objects
3. show that no such object exists

The method used is usually called *backtracking*. We will see how to reduce the search space if additional constraints are imposed on the paths. Finally, we show how careful optimization can lead to surprising algorithms for objects of a size where one would hardly expect to obtain a result at all. In fact, Gray cycles through the n-bit binary Lyndon words for all odd $n \leq 37$ are determined.

Terminology and conventions

We will use the terms *node* (instead of *vertex*) and *edge* (sometimes called *arc*). We restrict our attention to *directed graphs* (or *digraphs*) as undirected graphs are just the special case of these: an edge in an undirected graph corresponds to two antiparallel edges (think: 'arrows') in a directed graph.

A length-k *path* is a sequence of nodes where an edge leads from each node to its successor. A path is called *simple* if the nodes are pairwise distinct. We restrict our attention to simple paths of length N where N is the number of nodes of in the graph. We use the term *full path* for a simple path of length N.

If in a simple path there is an edge from the last node of the path to the starting node the path is a *cycle* (or *circuit*). A full path that is a cycle is called a *Hamiltonian cycle*, a graph containing such a cycle is called *Hamiltonian*.

We allow for *loops* (edges that start and point to the same node). Graphs that contain loops are called *pseudo graphs*. The algorithms used will effectively ignore loops. We disallow *multigraphs* (where multiple edges can start and end at the same two nodes), as these would lead to repeated output of identical objects.

The *neighbors* of a node are those nodes where outgoing edges point to. Neighbors can be reached with one step. The neighbors of a node a called *adjacent* to the node. The *adjacency matrix* of a graph with N nodes is a $N \times N$ matrix A where $A_{i,j} = 1$ if there is an edge from node i to node j, else $A_{i,j} = 0$.

While trivial to implement (and later modify) we will not use this kind of representation as the memory requirement would be prohibitive for large graphs.

19.1 Representation of digraphs

For our purposes a static implementation of the graph as arrays of nodes and (outgoing) edges will suffice. The container class `digraph` merely allocates memory for the nodes and edges. The correct initialization is left to the user [FXT: `class digraph` in graph/digraph.h]:

```
class digraph
{
public:
    ulong ng_;   // number of Nodes of Graph
    ulong *ep_;  // e[ep[k]], ..., e[ep[k+1]-1]: outgoing connections of node k
    ulong *e_;   // outgoing connections (Edges)
    ulong *vn_;  // optional: sorted values for nodes
    // if vn is used, then node k must correspond to vn[k]

public:
    digraph(ulong ng, ulong ne, ulong *&ep, ulong *&e, bool vnq=false)
        : ng_(0), ep_(0), e_(0), vn_(0)
    {
        ng_ = ng;
        ep_ = new ulong[ng_+1];
        e_ = new ulong[ne];
        ep = ep_;
        e = e_;
        if ( vnq )  vn_ = new ulong[ng_];
    }

    ~digraph()
    {
        delete [] ep_;
        delete [] e_;
        if ( vn_ )  delete [] vn_;
    }

 [--snip--]

    void get_edge_idx(ulong p, ulong &fe, ulong &en)  const
    // Setup fe and en so that the nodes reachable from p are
    //   e[fe], e[fe+1], ..., e[en-1].
    // Must have:  0<=p<ng
    {
        fe = ep_[p];   // (index of) First Edge
        en = ep_[p+1]; // (index of) first Edge of Next node
    }

 [--snip--]
    void print(const char *bla=0)  const;
};
```

The nodes reachable from node p could be listed using

```
// ulong p; // == position
cout << "The nodes reachable from node " << p << " are:" << endl;
ulong fe, en;
g_.get_edge_idx(p, fe, en);
for (ulong ep=fe; ep<en; ++ep)  cout << e_[ep] << endl;
```

Using our representation there is no cheap method to find the incoming edges. We will not need this information for our purposes. If the graph is known to be undirected, the same routine obviously lists the incoming edges.

Initialization routines for certain digraphs are declared in [FXT: graph/mk-special-digraphs.h]. A simple example is [FXT: `make_complete_digraph()` in graph/mk-complete-digraph.cc]:

```
digraph
make_complete_digraph(ulong n)
// Initialization for the complete graph.
```

```
4   {
5       ulong ng = n, ne = n*(n-1);
6
7       ulong *ep, *e;
8       digraph dg(ng, ne, ep, e);
9
10      ulong j = 0;
11      for (ulong k=0; k<ng; ++k)  // for all nodes
12      {
13          ep[k] = j;
14          for (ulong i=0; i<n; ++i)  // connect to all nodes
15          {
16              if ( k==i )  continue;  // skip loops
17              e[j++] = i;
18          }
19      }
20      ep[ng] = j;
21      return  dg;
22  }
```

We initialize the *complete graph* (the undirected graph that has edges between any two of its nodes) for $n = 5$ and print it [FXT: graph/graph-perm-demo.cc]:

```
digraph dg = make_complete_digraph(5);
dg.print("Graph =");
```

The output is

```
Graph =
Node: Edge0 Edge1 ...
  0:    1   2   3   4
  1:    0   2   3   4
  2:    0   1   3   4
  3:    0   1   2   4
  4:    0   1   2   3
  #nodes=5  #edges=20
```

For many purposes it suffices to implicitly represent the nodes as values p with $0 \le p < N$ where N is the number of nodes. If not, the values of the nodes have to be stored in the array `vn_[]`. One such example is a graph where the value of node p is the p-th (cyclically minimal) Lyndon word that we will meet at the end of this chapter. To make the search for a node by value reasonably fast the array `vn_[]` should be sorted so that binary search can be used.

19.2 Searching full paths

In order to search full paths starting from some position p_0 we need two additional arrays for the bookkeeping: A record `rv_[]` of the path so far, its k-th entry shall be p_k, the node visited at step k. Further a tag array `qq_[]` that shall contain a zero for nodes not visited so far, else one. The crucial parts of the implementation are [FXT: `class digraph_paths` in graph/digraph-paths.h]:

```
1   class digraph_paths
2   // Find all full paths in a directed graph.
3   {
4   public:
5       digraph &g_;  // the graph
6       ulong *rv_;  // Record of Visits: rv[k] == node visited at step k
7       ulong *qq_;  // qq[k] == whether node k has been visited yet
8    [--snip--]
9       // function to call with each path found with all_paths():
10      ulong (*pfunc_)(digraph_paths &);
11   [--snip--]
12      // function to impose condition with all_cond_paths():
13      bool (*cfunc_)(digraph_paths &, ulong ns);
14
15  public:
16      // graph/digraph.cc:
17      digraph_paths(digraph &g);
18      ~digraph_paths();
19   [--snip--]
20      bool path_is_cycle()  const;
```

```
21    [--snip--]
22       void print_path() const;
23    [--snip--]
24
25       // graph/digraphpaths-search.cc:
26       ulong all_paths(ulong (*pfunc)(digraph_paths &),
27                       ulong ns=0, ulong p=0, ulong maxnp=0);
28    private:
29       void next_path(ulong ns, ulong p);  // called by all_paths()
30    [--snip--]
31    };
```

We could have used a bit-array for the tag values `qq_[]`. It turns out that some additional information can be saved there as we will see in a moment.

To keep matters simple a recursive algorithm is used to search for (full) paths. The search is started via call to `all_paths()` [FXT: graph/digraph-paths.cc]:

```
1   ulong
2   digraph_paths::all_paths(ulong (*pfunc)(const digraph_paths &),
3                            ulong ns/*=0*/, ulong p/*=0*/, ulong maxnp/*=0*/)
4   // pfunc: function to visit (process) paths
5   // ns: start at node index ns (for fixing start of path)
6   // p: start at node value p   (for fixing start of path)
7   // maxnp: stop if maxnp paths were found
8   {
9       pct_ = 0;
10      cct_ = 0;
11      pfct_ = 0;
12      pfunc_ = pfunc;
13      pfdone_ = 0;
14      maxnp_ = maxnp;
15      next_path(ns, p);
16      return pfct_;  // Number of paths where pfunc() returned true
17  }
```

The search is done by the function `next_path()`:

```
1   void
2   digraph_paths::next_path(ulong ns, ulong p)
3   // ns+1 == how many nodes seen
4   // p == position (node we are on)
5   {
6       if ( pfdone_ )  return;
7
8       rv_[ns] = p;  // record position
9       ++ns;
10
11      if ( ns==ng_ )  // all nodes seen ?
12      {
13          pfunc_(*this);
14      }
15      else
16      {
17          qq_[p] = 1; // mark position as seen (else loops lead to errors)
18          ulong fe, en;
19          g_.get_edge_idx(p, fe, en);
20          ulong fct = 0;  // count free reachable nodes  // FCT
21          for (ulong ep=fe; ep<en; ++ep)
22          {
23              ulong t = g_.e_[ep]; // next node
24              if ( 0==qq_[t] )  // node free?
25              {
26                  ++fct;
27                  qq_[p] = fct; // mark position as seen: record turns  // FCT
28                  next_path(ns, t);
29              }
30          }
31          // if ( 0==fct )  { "dead end: this is a U-turn"; }  // FCT
32
33          qq_[p] = 0; // unmark position
34      }
35  }
```

The lines that are commented with `// FCT` record which among the free nodes is visited. The algorithm still works if these lines are commented out.

19.2.1 Paths in the complete graph: permutations

The program [FXT: graph/graph-perm-demo.cc] shows the paths in the complete graph from section 19.1, page 375. The version here is slightly simplified:

```
1   ulong pfunc_perm(digraph_paths &dp)
2   // Function to be called with each path:
3   //  print all but the first node.
4   {
5       const ulong *rv = dp.rv_;
6       ulong ng = dp.ng_;
7
8       cout << setw(4) << dp.pfct_ << ":  ";
9       for (ulong k=1; k<ng; ++k)  cout << " " << rv[k];
10      cout << endl;
11
12      return 1;
13  }
14
15  int
16  main(int argc, char **argv)
17  {
18      ulong n = 5;
19      digraph dg = make_complete_digraph(n);
20      digraph_paths dp(dg);
21
22      dg.print("Graph =");
23      cout << endl;
24
25      dp.all_paths(pfunc_perm, 0, 0, maxnp);
26      return 0;
27  }
```

The output is a listing of the permutations of the numbers $1, 2, 3, 4$ in lexicographic order (see section 10.1 on page 233).

```
    Graph =
    Node: Edge0 Edge1 ...
      0:    1   2   3   4
      1:    0   2   3   4
      2:    0   1   3   4
      3:    0   1   2   4
      4:    0   1   2   3
      #nodes=5  #edges=20

       0:   1 2 3 4
       1:   1 2 4 3
       2:   1 3 2 4
       3:   1 3 4 2
       4:   1 4 2 3
       5:   1 4 3 2
       6:   2 1 3 4
       7:   2 1 4 3
       8:   2 3 1 4
  [--snip--]
      21:   4 2 3 1
      22:   4 3 1 2
      23:   4 3 2 1
```

19.2.2 Paths in the De Bruijn graph: De Bruijn sequences

The graph with $2\,n$ nodes and two outgoing edges from node k to $2\,k \bmod 2\,n$ and $2\,k+1 \bmod 2\,n$ is called a *De Bruijn graph*. For $n = 8$ the graph is (printed horizontally):

```
    Node:  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
  Edge 0:  0  2  4  6  8 10 12 14  0  2  4  6  8 10 12 14
  Edge 1:  1  3  5  7  9 11 13 15  1  3  5  7  9 11 13 15
```

The graph has two loops at the first and the last node. All paths in the De Bruijn graph are cycles, the graph is Hamiltonian.

With n a power of two the paths correspond to the *De Bruijn sequences* (DBS) of length $2\,n$. The graph has as many full paths as there are DBSs and the zeros/ones in the DBS correspond to even/odd values of the nodes, respectively. This is demonstrated in [FXT: graph/graph-debruijn-demo.cc] (shortened):

```
ulong pq = 1; // whether and what to print with each cycle

ulong pfunc_db(digraph_paths &dp)
// Function to be called with each cycle.
{
    switch ( pq )
    {
    case 0:  break;  // just count
    case 1:  // print lowest bits (De Bruijn sequence)
        {
            ulong *rv = dp.rv_, ng = dp.ng_;
            for (ulong k=0; k<ng; ++k)  cout << (rv[k]&1UL ? '1' : '.');
            cout << endl;
            break;
        }
    [--snip--]
    }
    return 1;
}

int main(int argc, char **argv)
{
    ulong n = 8;
    NXARG(pq, "what to do in pfunc()");
    ulong maxnp = 0;
    NXARG(maxnp, "stop after maxnp paths (0: never stop)");
    ulong p0 = 0;
    NXARG(p0, "start position <2*n");

    digraph dg = make_debruijn_digraph(n);
    digraph_paths dp(dg);

    dg.print_horiz("Graph =");

    // call pfunc() with each cycle:
    dp.all_paths(pfunc_db, 0, p0, maxnp);

    cout << "n = " << n;
    cout << " (ng=" << dg.ng_ << ")";
    cout << "   #cycles = " << dp.cct_;
    cout << endl;

    return 0;
}

```

The macro `NXARG()` read one argument, it is defined in [FXT: nextarg.h]. The output is

```
arg 1: 8 == n  [size of graph == 2*n]  default=8
arg 2: 1 == pq  [what to do in pfunc()]  default=1
arg 3: 0 == maxnp  [stop after maxnp paths (0: never stop)]  default=0
arg 4: 0 == p0  [start position <2*n]  default=0
Graph =
  Node:  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
Edge 0:  0  2  4  6  8 10 12 14  0  2  4  6  8 10 12 14
Edge 1:  1  3  5  7  9 11 13 15  1  3  5  7  9 11 13 15
.1..11.1.1111...
.1..1111.1.11...
.1.1..11.1111...
.1.1..1111.11...
.1.11..1111.1...
.1.11.1...1111...
.1.1111...11.1...
.1.1111.1...11...
.11..1.1111.1...
.11.1..1.1111...
.11.1.1111..1...
.11.1111...1.1...
.1111..1.11.1...
.1111.1..1.11...
.1111.1.11..1...
.1111.11..1.1...
n = 8 (ng=16)   #cycles = 16
```

The algorithm is a very effective way of generating all DBSs of a given length, the 67,108,864 DBSs of length 64 are generated in 140 seconds when printing is disabled (set argument `pq` to zero), corresponding

to a rate of more than 450,000 DBSs per second.

Setting the argument `pq` to 4 prints the binary values of the successive nodes in the path horizontally:

```
./bin 32 4 1
    -#----##---#-#---###--#--#-##--##-#--####-#-#-###-##-######-----
    --#----##---#-#---###--#--#-##--##-#--####-#-#-###-##-######----
    ---#----##---#-#---###--#--#-##--##-#--####-#-#-###-##-######---
    ----#----##---#-#---###--#--#-##--##-#--####-#-#-###-##-######--
    -----#----##---#-#---###--#--#-##--##-#--####-#-#-###-##-######-
    ------#----##---#-#---###--#--#-##--##-#--####-#-#-###-##-######
```

The graph is constructed in a way that each word is the predecessor shifted by one with either zero or one inserted at position zero (top row).

The number of cycles in the De Bruijn graph equals the number of degree-n normal binary polynomials. An efficient procedure to compute it is given in section 40.6.3 on page 896. A closed form for the special case $n = 2^k$ is given in section 39.4 on page 862.

19.2.3 A modified De Bruijn graph: complement-shift sequences

A modification of the De Bruijn graph forces the nodes to be the complement of its predecessor shifted by one (again with either zero or one inserted at position zero). The routine to set up the graph is [FXT: `make_complement_shift_digraph()` in graph/mk-debruijn-digraph.cc]:

```
digraph
make_complement_shift_digraph(ulong n)
{
    ulong ng = 2*n, ne = 2*ng;
    ulong *ep, *e;
    digraph dg(ng, ne, ep, e);

    ulong j = 0;
    for (ulong k=0; k<ng; ++k)  // for all nodes
    {
        ep[k] = j;
        ulong r = (2*k) % ng;
        e[j++] = r;  // connect node k to node (2*k) mod ng
        r = (2*k+1) % ng;
        e[j++] = r;  // connect node k to node (2*k+1) mod ng
    }
    ep[ng] = j;
    // Here we have a De Bruijn graph.

    for (ulong k=0,j=ng-1; k<j; ++k,--j) swap2(e[ep[k]], e[ep[j]]); // end with ones
    for (ulong k=0,j=ng-1; k<j; ++k,--j) swap2(e[ep[k]+1], e[ep[j]+1]);

    return  dg;
}
```

The program [FXT: graph/graph-complementshift-demo.cc] gives:

```
arg 1: 32 == n  [size of graph == 2*n]  default=32
arg 2: 4 == pq  [whether and what to print with each path]  default=2
arg 3: 1 == maxnp  [stop after maxnp paths (0: never stop)]  default=1

------#---#-#-##----##--#--#-###---###-#--##-####--######-##-#-#
-######-###-#-#--####--##-##-#---###---#-##--#----##------#--#-#
-#------#---#-#-##----##--#--#-###---###-#--##-####--######-##-#
-#-######-###-#-#--####--##-##-#---###---#-##--#----##------#--#
-#-#------#---#-#-##----##--#--#-###---###-#--##-####--######-##
-#-#-######-###-#-#--####--##-##-#---###---#-##--#----##------#-
```

For n a power of two the sequence of binary words has the interesting property that the changes between successive words depend on their sequency: words with higher sequency change in less positions. Further, if two neighbor bits are set in some word then the next word never contains both bits again. Out of a run of $k > 0$ consecutive set bits in a word only one is contained in the next word.

See section 8.3 on page 207 for the connection with De Bruijn sequences.

19.3 Conditional search

Sometimes one wants to find paths that are subject to certain restrictions. Testing for each path found whether it has the desired property and discarding it if not is the most simple way. However, this will in many cases be extremely ineffective. An upper bound for the number of recursive calls of the search function `next_path()` with an graph with N nodes and an maximal number of v outgoing edges at each node is $u = N^v$.

For example, the graph corresponding to Gray codes of n-bit binary words has $N = 2^n$ nodes and (exactly) $c = n$ outgoing edges at each node. The graph is the n-dimensional hypercube.

n :	N	$u = N^c = N^n = 2^{n \cdot n}$
1:	2	2
2:	4	16
3:	8	512
4:	16	65,536
5:	32	33,554,432
6:	64	68,719,476,736
7:	128	562,949,953,421,312
8:	256	18,446,744,073,709,551,616
9:	512	2,417,851,639,229,258,349,412,352
10:	1024	1,267,650,600,228,229,401,496,703,205,376

To reduce the search space we use a function that rejects branches that would lead to a path not satisfying the imposed restrictions. A conditional search can be started via `all_cond_paths()` that has an additional function pointer `cfunc()` as argument that shall implement the condition. The corresponding method is declared as [FXT: graph/digraph-paths.h]:

```
bool (*cfunc_)(digraph_paths &, ulong ns);
```

Besides the data from the digraph-class it needs the number of nodes seen so far (`ns`) as an argument. A slight modification of the search routine does what we want [FXT: graph/search-digraph-cond.cc]:

```
1   void
2   digraph_paths::next_cond_path(ulong ns, ulong p)
3   {
4       [--snip--]  // same as next_path()
5       if ( ns==ng_ )  // all nodes seen ?
6       [--snip--]  // same as next_path()
7       else
8       {
9           qq_[p] = 1; // mark position as seen (else loops lead to errors)
10          ulong fe, en;
11          g_.get_edge_idx(p, fe, en);
12          ulong fct = 0;  // count free reachable nodes
13          for (ulong ep=fe; ep<en; ++ep)
14          {
15              ulong t = g_.e_[ep]; // next node
16              if ( 0==qq_[t] )  // node free?
17              {
18                  rv_[ns] = t;  // for cfunc()
19                  if ( cfunc_(*this, ns) )
20                  {
21                      ++fct;
22                      qq_[p] = fct; // mark position as seen: record turns
23                      next_cond_path(ns, t);
24                  }
25              }
26          }
27          qq_[p] = 0; // unmark position
28      }
29  }
```

The free node under consideration is written to the end of the record of visited nodes so `cfunc()` does not need it as an explicit argument.

19.3.1 Modular adjacent changes (MAC) Gray codes

```
 0:  .... 0    0    ...1  0          0:  .... 0    0    ...1  0
 1:  ...1 1    1    ..1.  1          1:  ...1 1    1    ..1.  1
 2:  ..11 2    3    .1..  2          2:  ..11 2    3    .1..  2
 3:  .111 3    7    1...  3          3:  .111 3    7    ..1.  1
 4:  1111 4   15    ...1  0          4:  .1.1 2    5    ...1  0
 5:  111. 3   14    ..1.  1          5:  .1.. 1    4    ..1.  1
 6:  11.. 2   12    ...1  0          6:  .11. 2    6    .1..  2
 7:  11.1 3   13    1...  3          7:  ..1. 1    2    1...  3
 8:  .1.1 2    5    ...1  0          8:  1.1. 2   10    .1..  2
 9:  .1.. 1    4    ..1.  1          9:  111. 3   14    ..1.  1
10:  .11. 2    6    .1..  2         10:  11.. 2   12    ...1  0
11:  ..1. 1    2    1...  3         11:  11.1 3   13    ..1.  1
12:  1.1. 2   10    ...1  0         12:  1111 4   15    .1..  2
13:  1.11 3   11    ..1.  1         13:  1.11 3   11    ..1.  1
14:  1..1 2    9    ...1  0         14:  1..1 2    9    ...1  0
15:  1... 1    8   [1...  3]       15:  1... 1    8   [1...  3]
```

Figure 19.3-A: Two four-bit modular adjacent changes (MAC) Gray codes. Both are cycles.

We search for Gray codes that have the *modular adjacent changes* (MAC) property: the values of successive elements of the delta sequence shall change by ± 1 modulo n. Two examples are show in figure 19.3-A. The sequence on the right side even has the stated property if the term 'modular' is omitted: It has the *adjacent changes* (AC) property.

As bit-wise cyclic shifts and reflections of MAC Gray codes are again MAC Gray codes we consider paths starting $0 \to 1 \to 2$ as canonical paths.

In the demo [FXT: graph/graph-macgray-demo.cc] the search is done as follows (shortened):

```
1   int main(int argc, char **argv)
2   {
3       ulong n = 5;
4       NXARG(n, "size in bits");
5       cf_nb = n;
6
7       digraph dg = make_gray_digraph(n, 0);
8       digraph_paths dp(dg);
9
10      ulong ns = 0, p = 0;
11      // MAC: canonical paths start as 0-->1-->3
12      {
13          dp.mark(0, ns);
14          dp.mark(1, ns);
15          p = 3;
16      }
17
18      dp.all_cond_paths(pfunc, cfunc_mac, ns, p, maxnp);
19      return 0;
20  }
```

The function used to impose the MAC condition is:

```
1   ulong cf_nb; // number of bits,  set in main()
2   bool cfunc_mac(digraph_paths &dp, ulong ns)
3   // Condition: difference of successive delta values (modulo n) == +-1
4   {
5       // path initialized, we have ns>=2
6       ulong p  = dp.rv_[ns],  p1 = dp.rv_[ns-1], p2 = dp.rv_[ns-2];
7       ulong c = p ^ p1,  c1 = p1 ^ p2;
8       if ( c & bit_rotate_left(c1,1,cf_nb) )  return  true;
9       if ( c1 & bit_rotate_left(c,1,cf_nb) )  return  true;
10      return false;
11  }
```

One finds paths for $n \leq 7$ ($n = 7$ takes about 15 minutes). Whether MAC Gray codes exist for $n \geq 8$ is unknown (none is found with a 40 hour search).

19.3.2 Adjacent changes (AC) Gray codes

For AC paths one can only discard track-reflected solutions, the canonical paths are those where the delta sequence starts with a values that is $\leq \lceil n/2 \rceil$. A function to impose the AC condition is

```
1   ulong cf_mt; // mid track < cf_mt,  set in main()
2   bool cfunc_ac(digraph_paths &dp, ulong ns)
3   // Condition: difference of successive delta values == +-1
4   {
5       if ( ns<2 )  return  (dp.rv_[1] < cf_mt); // avoid track-reflected solutions
6       ulong p  = dp.rv_[ns],  p1 = dp.rv_[ns-1], p2 = dp.rv_[ns-2];
7       ulong c = p ^ p1, c1 = p1 ^ p2;
8       if ( c & (c1<<1) )  return  true;
9       if ( c1 & (c<<1) )  return  true;
10      return false;
11  }
```

```
      0:  ..... 0    0    ..1..  2          0:  ..... 0    0    ....1  0
      1:  ..1.. 1    4    .1...  3          1:  ....1 1    1    ...1.  1
      2:  .11.. 2   12    1....  4          2:  ...11 2    3    ..1..  2
      3:  111.. 3   28    .1...  3          3:  ..111 3    7    .1...  3
      4:  1.1.. 2   20    ..1..  2          4:  .1111 4   15    ..1..  2
      5:  1.... 1   16    ...1.  1          5:  .1.11 3   11    ...1.  1
      6:  1..1. 2   18    ..1..  2          6:  .1..1 2    9    ..1..  2
      7:  1.11. 3   22    .1...  3          7:  .11.1 3   13    .1...  3
      8:  1111. 4   30    ..1..  2          8:  ..1.1 2    5    1....  4
      9:  11.1. 3   26    ...1.  1          9:  1.1.1 3   21    .1...  3
     10:  11... 2   24    ....1  0         10:  111.1 4   29    ..1..  2
     11:  11..1 3   25    ...1.  1         11:  11..1 3   25    ...1.  1
     12:  11.11 4   27    ..1..  2         12:  11.11 4   27    ..1..  2
     13:  11111 5   31    .1...  3         13:  11111 5   31    .1...  3
     14:  1.111 4   23    ..1..  2         14:  1.111 4   23    ..1..  2
     15:  1..11 3   19    ...1.  1         15:  1..11 3   19    ...1.  1
     16:  1...1 2   17    ..1..  2         16:  1...1 2   17    ....1  0
     17:  1.1.1 3   21    .1...  3         17:  1.... 1   16    ...1.  1
     18:  111.1 4   29    1....  4         18:  1..1. 2   18    ..1..  2
     19:  .11.1 3   13    .1...  3         19:  1.11. 3   22    .1...  3
     20:  ..1.1 2    5    ..1..  2         20:  1111. 4   30    ..1..  2
     21:  ....1 1    1    ...1.  1         21:  11.1. 3   26    ...1.  1
     22:  ...11 2    3    ..1..  2         22:  11... 2   24    ..1..  2
     23:  ..111 3    7    .1...  3         23:  111.. 3   28    .1...  3
     24:  .1111 4   15    ..1..  2         24:  1.1.. 2   20    1....  4
     25:  .1.11 3   11    ...1.  1         25:  ..1.. 1    4    .1...  3
     26:  .1..1 2    9    ....1  0         26:  .11.. 2   12    ..1..  2
     27:  .1... 1    8    ...1.  1         27:  .1... 1    8    ...1.  1
     28:  .1.1. 2   10    ..1..  2         28:  .1.1. 2   10    ..1..  2
     29:  .111. 3   14    .1...  3         29:  .111. 3   14    .1...  3
     30:  ..11. 2    6    ..1..  2         30:  ..11. 2    6    ..1..  2
     31:  ...1. 1    2   [...1.  1]       31:  ...1. 1    2   [...1.  1]
```

Figure 19.3-B: Two five-bit adjacent changes (AC) Gray codes that are cycles.

The program [FXT: graph/graph-acgray-demo.cc] allows searches for AC Gray codes. Two cycles for $n = 5$ are shown in figure 19.3-B. It turns out that such paths exist for $n \leq 6$ (the only path for $n = 6$ is shown in figure 19.3-C) but there is no AC Gray code for $n = 7$:

```
 time ./bin 7
arg 1: 7 == n  [size in bits]  default=5
arg 2: 0 == maxnp  [ stop after maxnp paths (0: never stop)]  default=0
n = 7   #pfct = 0
   #paths = 0   #cycles = 0
./bin 7  20.77s user 0.11s system 98% cpu 21.232 total
```

Nothing is known about the case $n \geq 8$.

Inspection of the AC Gray codes for different values of n result in a hand-woven algorithm. The function [FXT: `ac_gray_delta()` in comb/acgray.cc] computes the delta sequence for an AC Gray codes for $n \leq 6$:

```
1   void
2   ac_gray_delta(uchar *d, ulong ldn)
3   // Generate a delta sequence for an adjacent-changes (AC) Gray code
4   //  of length n=2**ldn where ldn<=6.
5   {
6       if ( ldn<=2 )  // standard Gray code
7       {
8           d[0] = 0;
```

```
 0:  ...... 0    0    ...1..  2       32:  1.1... 2   40    .1....  4
 1:  ...1.. 1    4    ....1.  1       33:  111... 3   56    ..1...  3
 2:  ...11. 2    6    ...1..  2       34:  11.... 2   48    ...1..  2
 3:  ....1. 1    2    ..1...  3       35:  11.1.. 3   52    ....1.  1
 4:  ..1.1. 2   10    ...1..  2       36:  11.11. 4   54    ...1..  2
 5:  ..111. 3   14    ....1.  1       37:  11..1. 3   50    ..1...  3
 6:  ..11.. 2   12    .....1  0       38:  111.1. 4   58    ...1..  2
 7:  ..11.1 3   13    ....1.  1       39:  11111. 5   62    ....1.  1
 8:  ..1111 4   15    ...1..  2       40:  1111.. 4   60    .....1  0
 9:  ..1.11 3   11    ..1...  3       41:  1111.1 5   61    ....1.  1
10:  ....11 2    3    ...1..  2       42:  111111 6   63    ...1..  2
11:  ...111 3    7    ....1.  1       43:  111.11 5   59    ..1...  3
12:  ...1.1 2    5    ...1..  2       44:  11..11 4   51    ...1..  2
13:  .....1 1    1    ..1...  3       45:  11.111 5   55    ....1.  1
14:  ..1..1 2    9    .1....  4       46:  11.1.1 4   53    ...1..  2
15:  .11..1 3   25    ..1...  3       47:  11...1 3   49    ..1...  3
16:  .1...1 2   17    ...1..  2       48:  111..1 4   57    .1....  4
17:  .1.1.1 3   21    ....1.  1       49:  1.1..1 3   41    ..1...  3
18:  .1.111 4   23    ...1..  2       50:  1....1 2   33    ...1..  2
19:  .1..11 3   19    ..1...  3       51:  1..1.1 3   37    ....1.  1
20:  .11.11 4   27    ...1..  2       52:  1..111 4   39    ...1..  2
21:  .11111 5   31    ....1.  1       53:  1...11 3   35    ..1...  3
22:  .111.1 4   29    .....1  0       54:  1.1.11 4   43    ...1..  2
23:  .111.. 3   28    ....1.  1       55:  1.1111 5   47    ....1.  1
24:  .1111. 4   30    ...1..  2       56:  1.11.1 4   45    .....1  0
25:  .11.1. 3   26    ..1...  3       57:  1.11.. 3   44    ....1.  1
26:  .1..1. 2   18    ...1..  2       58:  1.111. 4   46    ...1..  2
27:  .1.11. 3   22    ....1.  1       59:  1.1.1. 3   42    ..1...  3
28:  .1.1.. 2   20    ...1..  2       60:  1...1. 2   34    ...1..  2
29:  .1.... 1   16    ..1...  3       61:  1..11. 3   38    ....1.  1
30:  .11... 2   24    .1....  4       62:  1..1.. 2   36    ...1..  2
31:  ..1... 1    8    1.....  5       63:  1..... 1   32   [1.....  5]
```

Figure 19.3-C: The (essentially unique) AC Gray code for $n = 6$. While the path is a cycle in the graph the AC condition does not hold for the transition from the last to the first word.

```
9           if ( ldn==2 ) { d[1] = 1; d[2] = 0; }
10          return;
11      }
12
13      ac_gray_delta(d, ldn-1);  // recursion
14
15      ulong n = 1UL<<ldn;
16      ulong nh = n/2;
17      if ( 0==(ldn&1) )
18      {
19          if ( ldn>=6 )
20          {
21              reverse(d, nh-1);
22              for (ulong k=0;  k<nh;  ++k)  d[k] = (ldn-2) - d[k];
23          }
24
25          for (ulong k=0,j=n-2;  k<j;  ++k,--j)  d[j] = d[k];
26          d[nh-1] = ldn - 1;
27      }
28      else
29      {
30          for (ulong k=nh-2,j=nh-1;  0!=j;  --k,--j)  d[j] = d[k] + 1;
31          for (ulong k=2,j=n-2;  k<j;  ++k,--j)  d[j] = d[k];
32          d[0] = 0;
33          d[nh] = 0;
34      }
35  }
```

The Gray code is computed via

```
1   void
2   ac_gray(ulong *g, ulong ldn)
3   // Create an AC Gray code.
4   {
5       ulong n = 1UL<<ldn;
6       ALLOCA(uchar, d, n);
7       ac_gray_delta(d, ldn);
8       delta2gray(d, ldn, g);
9   }
```

where the routine `delta2gray()` is given in [FXT: comb/delta2gray.cc]:

```
1  void
2  delta2gray(const unsigned char *d, ulong ldn, ulong *g, ulong g0/*=0*/)
3  // Fill into g[0..N-1] (N=2**ldn) the Gray path
4  //  corresponding to the delta sequence d[0..N-2].
5  {
6      g[0] = g0;
7      ulong n = 1UL << ldn;
8      for (ulong k=0; k<n-1; ++k)  g[k+1] = g[k] ^ (1UL << d[k]);
9  }
```

The program [FXT: comb/acgray-demo.cc] can be used to create AC Gray codes for $n \leq 6$. For $n \geq 7$ the algorithm produces near-AC Gray codes, where the number of non-AC transitions equals $2^{n-5} - 1$ for odd values of n and $2^{n-5} - 2$ for n even:

```
 # non-AC transitions:
   n =0..6  #non-ac = 0
   n =  7   #non-ac = 3
   n =  8   #non-ac = 6
   n =  9   #non-ac = 15
   n = 10   #non-ac = 30
   n = 11   #non-ac = 63
   n = 12   #non-ac = 126
   ...
```

It seems likely that near-AC Gray codes with fewer non-AC transitions exist.

19.4 Edge sorting and lucky paths

The order of the nodes in the representation of the graph does not matter with finding paths as the algorithm does at no point refer to it. The order of the outgoing edges, however, *does* matter.

19.4.1 Edge sorting

Consider a large graph that has only a few paths. The calling tree of the recursive function `next_path()` obviously depends on the edge order. Thereby the first path can appear earlier or later in the search. 'Later' may well mean that the path is not found within any reasonable amount of time.

With a bit of luck one might find an ordering of the edges of the graph that will shorten the time span until the first path is found. The program [FXT: graph/graph-monotonicgray-demo.cc] searches for monotonic Gray codes and optionally sorts the edges of the graph. The method [FXT: `digraph::sort_edges()` in graph/digraph.cc] sorts the outgoing edges of each node according to a supplied comparison function.

The comparison function actually used imposes the lexicographic order shown in section 1.28 on page 71:

```
1  int my_cmp(const ulong &a, const ulong &b)
2  {
3      if ( a==b )  return  0;
4  #define CODE(x) lexrev2negidx(x);
5      ulong ca = CODE(a);
6      ulong cb = CODE(b);
7      return (ca<cb ? +1 : -1);
8  }
```

The choice was inspired by the observation that the bit-wise difference of successive elements in bit-lex order is either one or three. We search until the first path for 8-bit words is found: for the unsorted graph this task takes 1.14 seconds, for the sorted it takes 0.03 seconds.

19.4.2 Lucky paths

The first Gray code found in the hypercube graph with randomized edge order is shown in figure 19.4-A (left). The corresponding path (as reported by [FXT: `digraph_paths::print_turns()` in graph/digraph-paths.cc]) is described in the right column. Here `nn` is the number of neighbors of `node`, `xe` is the index

```
 0:  .... 0    0    1...  3          step: node -> next [xf xe / nn]
 1:  1... 1    8    ..1.  1             0:    0 ->   8  [ 0  0 /  4]
 2:  1.1. 2   10    .1..  2             1:    8 ->  10  [ 0  0 /  4]
 3:  111. 3   14    ...1  0             2:   10 ->  14  [ 0  0 /  4]
 4:  1111 4   15    1...  3             3:   14 ->  15  [ 0  0 /  4]
 5:  .111 3    7    .1..  2             4:   15 ->   7  [ 0  0 /  4]
 6:  ..11 2    3    ...1  0             5:    7 ->   3  [ 0  1 /  4]
 7:  ..1. 1    2    .1..  2             6:    3 ->   2  [ 1  2 /  4]
 8:  .11. 2    6    ..1.  1             7:    2 ->   6  [ 0  3 /  4]
 9:  .1.. 1    4    1...  3             8:    6 ->   4  [ 0  0 /  4]
10:  11.. 2   12    ...1  0             9:    4 ->  12  [ 1  3 /  4]
11:  11.1 3   13    1...  3            10:   12 ->  13  [ 0  0 /  4]
12:  .1.1 2    5    .1..  2            11:   13 ->   5  [ 0  1 /  4]
13:  ...1 1    1    1...  3            12:    5 ->   1  [ 0  3 /  4]
14:  1..1 2    9    ..1.  1            13:    1 ->   9  [ 0  2 /  4]
15:  1.11 3   11   [1.11  -]           14:    9 ->  11  [ 0  0 /  4]
                                    Path: #non-first-free turns = 2
```

Figure 19.4-A: A Gray code in the hypercube graph with randomized edge order (left) and the path description (right, see text).

---

of the neighbor (`next`) in the list of edges of `node`. Finally `xf` is the index among the *free* nodes in the list. The latter corresponds to the value `fct-1` in the function `next_path()` given on page 376.

If `xf` equals zero at some step then the first free neighbor was visited next. If `xf` is nonzero then a dead end was reached in the course of the search and there was at least one U-turn. If the path is not the first found then the U-turn might well correspond to a previous path.

If there was no U-turn then the number of *non-first-free* turns is zero (the number is given as the last line of the report). If it is zero we call the path found a *lucky path*. For each given ordering of the edges and each starting position of the search there is at most one lucky path and if there is, it is the first path found.

When the first path is a lucky path then the search effectively 'falls through': the number of operations is a constant times the number of edges. That is, if a lucky path exists it is found almost immediately even for huge graphs.

19.5 Gray codes for Lyndon words

We search Gray codes for n-bit binary Lyndon words where n is a prime. Here is a Gray code for the 5-bit Lyndon words that is a cycle:

```
...1
..11
.111
.1111
..111
..1.1
```

An important application of such Gray codes is the construction of so-called *single track Gray codes* which can be obtained by appending rotated versions of the block. The following is a single track Gray code based on the block given. At each stage, the block is rotated by two positions (horizontal format):

```
######  --##--  -####-  ------  ---###
-####-  ------  ---###  ######  --##--
---###  ######  --##--  -####-  ------
--##--  -####-  ------  ---###  ######
------  ---###  ######  --##--  -####-
```

The *transition count* (the number of zero-one transitions) is by construction the same for each track. The all-zero and the all-one words are missing in the Gray code, its length equals $2^n - 2$.

```
 0:   .....1    .....1    .....1    .....1    .....1    .....1
 1:   ....11    ...1..1   ...1..1   ...1..1   ...1..1   ....11
 2:   ...111    ...11.1   ...11.1   ...11.1   ...1.11   ....111
 3:   ...1111   ..111.1   ..111.1   ....1.1   .1.1.11   ....1.1
 4:   ..11111   ..1.1.1   ..11111   ..1.1.1   .1.1111   ..1.1.1
 5:   ..11.11   ..1.111   .111111   ..111.1   .111111   ..111.1
 6:   ..1..11   .11.111   .11.111   ..11111   .11.111   ...11.1
 7:   ..1.111   .111111   ..1.111   ..11.11   ..1.111   ...1..1
 8:   .11.111   .1.1111   ..1.1.1   ..1..11   ..1.1.1   ...1.11
 9:   .111111   .1.1.11   ....1.1   ..1.111   ..111.1   ...1111
10:   .1.1111   ...1.11   ....111   .11.111   ...11.1   ..11111
11:   .1.1.11   ...1111   ...1111   .111111   ....1.1   ..11.11
12:   ...1.11   ..11111   .1.1111   .1.1111   ....111   ..1..11
13:   ...1..1   ..11.11   .1.1.11   .1.1.11   .....11   ..1.111
14:   ...11.1   ..1..11   ...1.11   ...1.11   ..1..11   .11.111
15:   ..111.1   .....11   ..11.11   ...1111   ..11.11   .111111
16:   ..1.1.1   ....111   ..1..11   ....111   ..11111   .1.1111
17:   ....1.1   ....1.1   .....11   .....11   ...1111   .1.1.11
```

Figure 19.5-A: Various Gray codes through the length-7 binary Lyndon words. The first four are cycles.

19.5.1 Graph search with edge sorting

Gray codes for the 7-bit binary Lyndon words like those shown in figure 19.5-A can easily be found by a graph search. In fact, all of them can be generated in short time: for $n = 7$ there are 395 Gray codes (starting with the word `0000..001`) of which 112 are cycles.

```
 k :  [ node]  lyn_dec lyn_bin #rot rot(lyn)     diff  delta
 0 :  [    0]        1  .....1  0   .....1   .....1  0
 1 :  [    1]        3  ....11  0   ....11   ....1.  1
 2 :  [    3]        7  ...111  0   ...111   ...1..  2
 3 :  [    7]       15  ..1111  0   ..1111   ..1...  3
 4 :  [   13]       31  .11111  0   .11111   .1....  4
 5 :  [   17]       63  .11111  0   .11111   .1....  5
 6 :  [   15]       47  .1.1111 0   .1.1111  ..1...  4
 7 :  [   10]       23  ..1.111 1   .1.111.  .....1  0
 8 :  [   16]       55  .11.111 1   11.111.  1.....  6
 9 :  [   11]       27  ..11.11 2   11.11..  ....1.  1
10 :  [    5]       11  ...1.11 2   .1.11..  1.....  6
11 :  [   14]       43  .1.1.11 2   .1.11.1  .....1  0
12 :  [    6]       13  ...11.1 0   ...11.1  .1....  5
13 :  [   12]       29  ..111.1 0   ..111.1  ..1...  4
14 :  [    8]       19  ..1..11 3   ..11..1  ...1..  2
15 :  [    4]        9  ...1..1 0   ...1..1  ..1...  4
16 :  [    9]       21  ..1.1.1 3   .1.1..1  .1....  5
17 :  [    2]        5  ....1.1 3   .1.1...  .....1  0
```

Figure 19.5-B: A Gray code through the length-7 binary Lyndon words found by searching through all 7-bit binary words.

The search for such a path for the next prime, $n = 11$, does not seem to give a result in reasonable time. If we do not insist on a Gray code through the cyclic minima but arbitrary rotations of the Lyndon words then more Gray codes exist. For that purpose nodes are declared adjacent if there is any cyclic rotation of the second node's value that differs in exactly one bit to the first node's value. The cyclic rotations can be recovered easily after a path is found. This is done in [FXT: graph/graph-lyndon-gray-demo.cc] whose output is shown in figure 19.5-B. Still, already for $n = 11$ we do not get a result. As the corresponding graph has 186 nodes and 1954 edges, this is not a surprise.

Now we try edge sorting, we sort the edges according to the comparison function [FXT: graph/lyndon-cmp.cc]

```
1   int lyndon_cmp0(const ulong &a, const ulong &b)
2   {
3       int bc = bit_count_cmp(a, b);
4       if ( bc )  return  -bc;  // more bits first
5       else
6       {
7           if ( a==b )  return 0;
8           return  (a>b ?  +1 : -1);  // bigger numbers last
9       }
```

```
10  }
```

where `bit_count_cmp()` is defined in [FXT: bits/bitcount.h]:

```
1   static inline int bit_count_cmp(const ulong &a, const ulong &b)
2   {
3       ulong ca = bit_count(a);
4       ulong cb = bit_count(b);
5       return  ( ca==cb ? 0 : (ca>cb ? +1 : -1) );
6   }
```

```
     k :  [ node]  lyn_dec       lyn_bin #rot      rot(lyn)           diff  delta
     0 :  [    0]         1  ...........1  0   ...........1   ...........1  0
     1 :  [    1]         3  ..........11  0   ..........11   ..........1.  1
     2 :  [    3]         7  .........111  0   .........111   .........1..  2
     3 :  [    7]        15  ........1111  0   ........1111   ........1...  3
     4 :  [   15]        31  .......11111  0   .......11111   .......1....  4
     5 :  [   31]        63  ......111111  0   ......111111   ......1.....  5
     6 :  [   63]       127  .....1111111  0   .....1111111   .....1......  6
     7 :  [  125]       255  ....11111111  0   ....11111111   ....1.......  7
     8 :  [  239]       511  ...111111111  0   ...111111111   ...1........  8
     9 :  [  417]      1023  ..1111111111  0   ..1111111111   ..1.........  9
    10 :  [  589]      2047  .11111111111  0   .11111111111   .1.......... 10
    11 :  [  629]      4095  .11111111111  0   .11111111111   .1.......... 11
    12 :  [  618]      3071  .1.1111111111  0   .1.1111111111  ..1......... 10
    13 :  [  514]      1535  ..1.111111111  1   .1.111111111.  ...........1  0
    14 :  [  624]      3583  .11.111111111  1   11.111111111.  1........... 12
    15 :  [  550]      1791  ..11.11111111  2   11.11111111..  ..........1.  1
    16 :  [  626]      3839  .111.11111111  2   11.11111111.1  ...........1  0
    17 :  [  567]      1919  ..111.1111111  3   11.1111111..1  .........1..  2
    18 :  [  627]      3967  .1111.1111111  3   11.1111111.11  ..........1.  1
    19 :  [  576]      1983  ..1111.111111  4   11.111111..11  ........1...  3
    20 :  [  628]      4031  .11111.111111  4   11.111111.111  .........1..  2
    21 :  [  581]      2015  ..11111.11111  5   11.11111..111  .......1....  4
    22 :  [  404]       991  ...1111.11111  5   11.11111...11  .........1..  2
    23 :  [  614]      3039  .1.1111.11111  5   11.11111.1.11  ........1...  3
    24 :  [  508]      1519  ..1.1111.1111  6   11.1111..1.11  ......1.....  5
    25 :  [  584]      2031  ..111111.1111  6   11.1111..1111  .........1..  2
     [--snip--]
   615 :  [    4]         9  ........1..1  5   ....1..1.....  ..1......... 10
   616 :  [   36]        73  ......1..1..1  2   ....1..1..1..  .........1..  2
   617 :  [   32]        65  ......1.....1  2   ....1.....1..  .......1.....  5
   618 :  [   33]        67  ......1....11  2   ....1....11..  ........1...  3
   619 :  [  153]       323  ....1.1....11  2   ..1.1....11..  ..1......... 10
   620 :  [   65]       133  .....1....1.1  8   ..1.1.....1..  .........1...  3
   621 :  [  154]       325  ....1.1...1.1  2   ..1.1...1.1..  ........1....  4
   622 :  [   79]       161  .....1.1....1 10   ..1.....1.1..  ....1........  8
   623 :  [   16]        33  .......1....1 10   ..1.......1..  ........1....  4
   624 :  [  126]       265  ....1....1..1  2   ..1....1..1..  .......1.....  5
   625 :  [  145]       305  ....1..11...1 10   ..1....1..11.  ..........1.  1
   626 :  [  130]       273  ....1...1...1 10   ..1....1...1.  .........1..  2
   627 :  [  188]       401  ....11..1...1 10   ..1....11..1.  ........1....  4
   628 :  [   71]       145  .....1..1...1 10   ..1.....1..1.  .......1.....  5
   629 :  [    8]        17  ........1...1 10   ..1........1.  ........1....  4
```

Figure 19.5-C: Begin and end of a Gray cycle through the 13-bit binary Lyndon words.

We find a Gray code (which also is a cycle) for $n = 11$ immediately. Same for $n = 13$, again a cycle. The graph for $n = 13$ has 630 nodes and 8,056 edges, so finding a path is quite unexpected. The cycle found starts and ends as shown in figure 19.5-C.

For next candidate ($n = 17$) we do no find a Gray code within many hours of search. No surprise for a graph with 7,710 nodes and 130,828 edges.

We try another edge sorting scheme, an ordering based on the binary Gray code [FXT: graph/lyndon-cmp.cc]:

```
1   int lyndon_cmp2(const ulong &a, const ulong &b)
2   {
3       if ( a==b )  return 0;
4   #define CODE(x) gray_code(x)
5       ulong ta = CODE(a), tb = CODE(b);
6       return  ( ta<tb ? +1 : -1);
7   }
```

There we go, we find a cycle for $n = 17$ and all smaller primes. All are cycles and all paths are lucky paths. The following edge sorting scheme also leads to Gray codes for all prime n where $3 \le n \le 17$:

```
1  int lyndon_cmp3(const ulong &a, const ulong &b)
2  {
3      if ( a==b )  return 0;
4  #define CODE(x) inverse_gray_code(x)
5      ulong ta = CODE(a), tb = CODE(b);
6      return  ( ta<tb ? +1 : -1);
7  }
```

Same for $n = 19$, the graph has 27,594 nodes and 523,978 edges. Indeed the sorting scheme leads to cycles for all odd $n \le 27$!. All these paths are lucky paths, a fact that we can exploit for an optimized search.

19.5.2 An optimized algorithm

n	number of nodes	tag-size	time
23	364,722	0.25 MB	1 sec
25	1,342,182	1 MB	3 sec
27	4,971,066	4 MB	12 sec
29	18,512,790	16 MB	1 min
31	69,273,666	64 MB	4 min
33	260,301,174	256 MB	16 min

n	number of nodes	tag-size	time
35	981,706,830	1 GB	1 h
37	3,714,566,310	4 GB	7 h
39	14,096,303,342	16 GB	2 d
41	53,634,713,550	64 GB	10 d
43	204,560,302,842	256 GB	>40 d
45	781,874,934,568	1 TB	>160 d

Figure 19.5-D: Memory and (approximate) time needed for computing Gray codes with n-bit Lyndon words. The number of nodes equals the number of length-n necklaces minus two. The size of the tag array equals $2^n/4$ bits or $2^n/32$ bytes.

With edge sorting functions that lead to a lucky path we can discard most of the data used with graph searching. We only need to keep track of whether a node has been visited so far. A tag-array ([FXT: ds/bitarray.h], see section 4.6 on page 160) suffices.

With n-bit Lyndon words the amount of tag-bits needed is 2^n. Find an implementation of the algorithm as [FXT: `class lyndon_gray` in graph/lyndon-gray.h].

If only the cyclical minima of the values are tagged then only $2^n/2$ bits are needed if the access to the single necklace consisting of all ones is treated separately. This variant of the algorithm is activated by uncommenting the line `#define ALT_ALGORITM`. Noting that the lowest bit in a necklace is always one we need only $2^n/4$ bits: simply shift the words right by one before testing or writing to the tag array. This can be achieved by additionally uncommenting the line `#define ALTALT` in the file.

When a node is visited the algorithm creates a table of neighbors and selects the minimum among the free nodes with respect to the edge sorting function used. The table of neighbors is discarded then in order to minimize memory usage.

When no neighbor is found the number of nodes visited so far is returned. If this number equals the number of n-bit Lyndon words a lucky path was found. With composite n a Gray code for n-bit necklaces with the exception of the all-ones and the all-zeros word will be searched.

Four flavors of the algorithm have been found so far, corresponding to edge sorting with the 3rd, 5th, 21th, and 29th power of the Gray code. We refer to these functions as comparison functions 0, 1, 2, and 3, respectively. All of these lead to cycles for all primes $n \le 31$. The resources needed with larger values of n are shown in figure 19.5-D.

Using a 64-bit machine equipped with more than 4 Gigabyte of RAM it can be verified that three of the edge sorting functions lead to a Gray cycles also for $n = 37$, the 3rd power version fails. One of the

sorting functions *may* lead to a Gray code for $n = 41$.

```
% ./bin 7 2 0  # 7 bits, full output, comparison function 0
n = 7   #lyn = 18
   1:    .....1   0  .....1   .....1  0
   2:    ...1..1  0  ...1..1  ...1...  3
   3:    ..1..11  3  ..11..1  ..1....  4
   4:    ..1.111  3  .111..1  .1.....  5
   5:    .1.1111  2  .1111.1  ....1..  2
   6:    .1.1.11  2  .1.11.1  ..1....  4
   7:    .11.111  5  11.11.1  1......  6
   8:    .111111  2  11111.1  ..1....  4
   9:    ..11111  2  11111..  .....1  0
  10:    ..111.1  2  111.1..  ...1...  3
  11:    ..1.1.1  2  1.1.1..  .1.....  5
  12:    ....1.1  2  ..1.1..  1......  6
  13:    ...1.11  1  ..1.11.  .....1.  1
  14:    ..11.11  1  .11.11.  .1.....  5
  15:    ...11.1  2  .11.1..  ....1.  1
  16:    ...1111  2  .1111..  ...1...  3
  17:    ....111  2  ..111..  .1.....  5
  18:    .....11  2  ...11..  ..1....  4
last = .....11     crc=0b14a5846c41d57f
n = 7   #lyn = 18   #= 18
```

Figure 19.5-E: A Gray code for 7-bit Lyndon words.

A program to compute the Gray codes is [FXT: graph/lyndon-gray-demo.cc], four arguments can be given:

```
arg 1: 13 == n  [ a prime < BITS_PER_LONG ]  default=17
arg 2: 1 == wh  [printing: 0==>none, 1==>delta seq., 2==>full output]  default=1
arg 3: 3 == ncmp  [use comparison function (0,1,2,3)]  default=2
arg 4: 0 == testall  [special: test all odd values <= value]  default=0
```

An example with full output is given in figure 19.5-E. A 64-bit CRC (see section 1.32 on page 84) is computed from the delta sequence (rightmost column) and printed with the last word.

```
% ./bin 13 1 2  # 13 bits, delta seq. output, comparison function 2
 n = 13   #lyn = 630
 06B57458354645962546436734A74684A106C0145120825747A745247AC8564567018A7654647484A756A546457CA1ACBC1C
 856BA9A64B974565486456596452194259425215315BC82BC75BA02926256354267A462475A3ACB9761560C37412583758CA5624
 B8C6A6C6A87A9C20CBA45340420145405231290756976515631602042 30A7BA31C1485C6105201510490BCA891BA9B1B9AC0
 A9A89B898A565B878574586574784 5A9546702305A41275315458767465747A8457845470379A8586B0A76985787679767 59
 A976567686A5676 56A576B86581305A20AB0ACB0AB535234382354653252475 63A432532A372354657643572373624634642
 4532397423435235653236423263235234327532342325396926853234232582642436823632346236582342383242327
 5232423253234326423242 35323423
 last = ..........11     crc=568dab04b55aa2fb
 n = 13   #lyn = 630   #= 630


% ./bin 13 1 3  # 13 bits, delta seq. output, comparison function 3
n = 13   #lyn = 630
  06B57458354645962546436735371CA8B1587BA7610635285A0C2484B9713476B689A897AC98768968B9A106326016261050
  1424B8979A78987B97898C9892194131531369831428 1687BCB9469C489C6210205B050A1A7A4568A9BC5CB79AB647B74812
  0AB30BC1A131ACB120B0164CA1CABA121ABACA2B0BACAB184578678498958486764 6A84561916546947457875458654901 37
  4020103101210427017121650745 7B854606C16BC52380136516413016 4BC7987A09872CBA9A87A20B787AC9B7CBA834C0C1
  3C341C1042010C14C01C4145878546 45A854C95035A6A9570A9756586B9B5969580A0872C3123B0CB316BC6C0B21B2C0C2C0
  5301C0530CB1C1530C01CB0BC20CBC0CB1C87565756865A75A65A40898A898B91CA898A8B898A81BC8A9ACA989AB817A9BC1
  BA9ABA9CA9AB918A1CACBAC9BCB0BC
  last = ..........11     crc=745def277b1fbed0
  n = 13   #lyn = 630   #= 630
```

Figure 19.5-F: Delta sequences for two different Gray codes for 13-bit Lyndon words.

For larger n one might want to print only the delta sequence, this is shown in figure 19.5-F. The CRC allows us to conveniently determine whether two delta sequences are different. Different sequences sometimes start identically.

For still larger values of n even the delta sequence tends to get huge (for example, with $n = 37$ the sequence would be approximately 3.7 GB). One can suppress all output except for a progress indication, as shown in figure 19.5-G. Here the CRC checksum is updated only with every (cyclically unadjusted) 2^{16}-th Lyndon word.

Sometimes a Gray code through the necklaces (except for the all-zeros and all-ones words) is also found for composite n. Comparison functions 0, 1, and 2 lead to Gray codes (which are cycles) for all odd

```
% ./bin 29 0 0  # 29 bits, output=progress, comparison function 0
n = 29   #lyn = 18512790
.............    1048576 ( 5.66406 % )   crc=ceabc5f2056be699
.............    2097152 ( 11.3281 % )   crc=76dd94f1a554b50d
.............    3145728 ( 16.9922 % )   crc=6b39957f1e141f4d
.............    4194304 ( 22.6563 % )   crc=53419af1f1185dc0
.............    5242880 ( 28.3203 % )   crc=45d45b193f8ee566
.............    6291456 ( 33.9844 % )   crc=95a24c824f56e196
.............    7340032 ( 39.6484 % )   crc=003ee5af5b248e34
.............    8388608 ( 45.3125 % )   crc=23cb74d3ea0c4587
.............    9437184 ( 50.9766 % )   crc=896fd04c87dd0d43
.............   10485760 ( 56.6406 % )   crc=b00d8c899f0fc791
.............   11534336 ( 62.3047 % )   crc=d148f1b95b23eeab
.............   12582912 ( 67.9688 % )   crc=82971e2ed4863050
.............   13631488 ( 73.6328 % )   crc=f249ad5b4fed252d
.............   14680064 ( 79.2969 % )   crc=909821d0c7246a98
.............   15728640 ( 84.9609 % )   crc=1c5d68e38e55b3ca
.............   16777216 ( 90.625 % )    crc=0e64f82c67c79cf1
.............   17825792 ( 96.2891 % )   crc=62c17b9f3c644396
.........
last = ............................11     crc=5736fc9365da927e
n = 29   #lyn = 18512790   #= 18512790
```

Figure 19.5-G: Computation of a Gray code through the 29-bit Lyndon words. Most output is suppressed, only the CRC is printed at certain checkpoints.

$n \leq 33$. Gray cycles are also found with comparison function 3, except for $n = 21$, 27, and 33. All functions give Gray cycles also for $n = 4$ and $n = 6$. The values of n for which no Gray code could be found are the even values ≥ 8.

19.5.3 No Gray codes for even $n \geq 8$

As the parity of the words in a Gray code sequence alternates between one and zero the difference between the numbers words of odd and even weight must be zero or one. If it is ones, no Gray cycle can exist because the parity of the first and last word is identical. We use the relations from section 17.3 on page 362.

For Lyndon words of odd length there are the same number of words for odd and even weight by symmetry, see figure 17.3-B on page 363. So a Gray code (and also a Gray cycle) can exist.

For even lengths the sequence of numbers of Lyndon words of odd and even weights start as:

```
   n:   2, 4, 6,  8, 10,  12,  14,   16,   18,    20,    22,     24,      26, ...
 odd:   1, 2, 5, 16, 51, 170, 585, 2048, 7280, 26214, 95325, 349520, 1290555, ...
even:   0, 1, 4, 14, 48, 165, 576, 2032, 7252, 26163, 95232, 349350, 1290240, ...
diff:   1, 1, 1,  2,  3,   5,   9,   16,   28,    51,    93,    170,     315, ...
```

The last row gives the differences, entry A000048 of [240]. All entries for $n \geq 8$ are greater than one, so no Gray code does exist.

For the number of necklaces we obtain, for $n = 2, 4, 6, \ldots$

```
   n:   2, 4, 6,  8, 10,  12,  14,   16,   18,    20,    22,     24,      26, ...
 odd:   1, 2, 6, 16, 52, 172, 586, 2048, 7286, 26216, 95326, 349536, 1290556, ...
even:   2, 4, 8, 20, 56, 180, 596, 2068, 7316, 26272, 95420, 349716, 1290872, ...
diff:   1, 2, 2,  4,  4,   8,  10,   20,   30,    56,    94,    180,     316, ...
```

The (absolute) difference of both sequences is entry A000013 of [240]. We see that for $n \geq 4$ the numbers are greater than one, so no Gray code exists.

If we exclude the all-ones and all-zeros words, then the differences are

```
   n:   2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22,  24,  26, ...
diff:   1, 0, 0, 2,  2,  6,  8, 18, 28, 54, 92, 178, 314, ...
```

And again, no Gray code exists for $n \geq 8$. That is, we have found Gray codes, and even cycles, for all computational feasible sizes were they can exist.

Part III

Fast orthogonal transforms

Chapter 20

The Fourier transform

We introduce the discrete Fourier transform and give algorithms for its fast computation. Implementations and optimization considerations for complex and real valued transforms are given. The fast Fourier transforms (FFTs) are the basis of the algorithms for fast convolution (given in chapter 21) which are in turn the core of the fast high precision multiplication routines described in chapter 27. The number theoretic Fourier transforms (NTTs) are treated in chapter 25. Algorithms for Fourier transforms based on fast convolution are given in chapter 21, the transform for arbitrary length (Bluestein's algorithm) in section 21.5, and prime-length transforms (Rader's algorithm) in section 21.6.

20.1 The discrete Fourier transform

The *discrete Fourier transform* (DFT or simply FT) of a complex sequence $a = [a_0, a_1, \ldots, a_{n-1}]$ of length n is the complex sequence $c = [c_0, c_1, \ldots, c_{n-1}]$ defined by

$$c = \mathcal{F}\big[\, a \,\big] \tag{20.1-1a}$$

$$c_k := \frac{1}{\sqrt{n}} \sum_{x=0}^{n-1} a_x \, z^{+x\,k} \quad \text{where} \quad z = e^{2\,\pi\,i/n} \tag{20.1-1b}$$

z is a primitive n-th root of unity: $z^n = 1$ and $z^j \neq 1$ for $0 < j < n$.

Back-transform (or *inverse discrete Fourier transform*, IDFT or simply IFT) is then

$$a = \mathcal{F}^{-1}\big[\, c \,\big] \tag{20.1-2a}$$

$$a_x := \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} c_k \, z^{-x\,k} \tag{20.1-2b}$$

To see this, consider element y of the IFT of the FT of a:

$$\mathcal{F}^{-1}\big[\, \mathcal{F}\big[\, a \,\big] \,\big]_y = \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} \frac{1}{\sqrt{n}} \sum_{x=0}^{n-1} (a_x \, z^{x\,k}) \, z^{-y\,k} \tag{20.1-3a}$$

$$= \frac{1}{n} \sum_x a_x \sum_k (z^{x-y})^k \tag{20.1-3b}$$

Now $\sum_k (z^{x-y})^k = n$ for $x = y$ and zero else. This is because z is an n-th primitive root of unity: with $x = y$ the sum consists of n times $z^0 = 1$, with $x \neq y$ the summands lie on the unit circle (on the vertices

of an equilateral polygon with center zero) and add up to zero. Therefore the whole expression is equal to

$$\frac{1}{n}\, n \sum_x a_x\, \delta_{x,y} \quad = \quad a_y \quad \text{where} \quad \delta_{x,y} := \begin{cases} 1 & (x = y) \\ 0 & (x \neq y) \end{cases} \tag{20.1-4a}$$

Here we will call the FT with the plus in the exponent the forward transform. The choice is actually arbitrary, engineers seem to prefer the minus for the forward transform, mathematicians the plus. The sign in the exponent is called the *sign of the transform*.

The FT is a linear transform. That is, for $\alpha, \beta \in \mathbb{C}$

$$\mathcal{F}\big[\,\alpha\, a + \beta\, b\,\big] \quad = \quad \alpha\, \mathcal{F}\big[\,a\,\big] + \beta\, \mathcal{F}\big[\,b\,\big] \tag{20.1-5}$$

Further *Parseval's equation* holds, the sum of squares of the absolute values is identical for a sequence and its Fourier transform:

$$\sum_{x=0}^{n-1} |a_x|^2 \quad = \quad \sum_{k=0}^{n-1} |c_k|^2 \tag{20.1-6}$$

A straightforward implementation of the discrete Fourier transform, that is, the computation of n sums each of length n, requires $\sim n^2$ operations.

```
1   void slow_ft(Complex *f, long n, int is)
2   {
3       Complex h[n];
4       const double ph0 = is*2.0*M_PI/n;
5       for (long w=0; w<n; ++w)
6       {
7           Complex t = 0.0;
8           for (long k=0; k<n; ++k)
9           {
10              t +=  f[k] * SinCos(ph0*k*w);
11          }
12          h[w] = t;
13      }
14      copy(h, f, n);
15  }
```

This is [FXT: `slow_ft()` in fft/slowft.cc]. The variable `is`$= \pm 1$ is the sign of the transform, the function `SinCos(x)` returns a `Complex(cos(x), sin(x))`.

Note that the normalization factor $\frac{1}{\sqrt{n}}$ in front of the FT sums has been left out. The inverse of the transform with one sign is the transform with the opposite sign followed by a multiplication of each element by $\frac{1}{n}$. One has to keep in mind if that the sum of squares of the original sequence and its transform are equal up to a factor $1/\sqrt{n}$.

A *fast Fourier transform* (FFT) algorithm is an algorithm that improves the operation count to proportional $n \sum_{k=1}^{m}(p_k - 1)$, where $n = p_1 p_2 \cdots p_m$ is a factorization of n. In case of a power $n = p^m$ the value computes to $n\,(p-1)\,\log_p(n)$. In the special case $p = 2$ even $n/2\,\log_2(n)$ (complex) multiplications suffice. There are several different FFT algorithms with many variants.

20.2 Summary of definitions of Fourier transforms *

We summarize the definitions of the continuous, semi-continuous and the discrete Fourier transform.

The continuous Fourier transform

The (continuous) *Fourier transform* (FT) of a function $f : \mathbb{C}^n \to \mathbb{C}^n, \quad \vec{x} \mapsto f(\vec{x})$ is defined by

$$F(\vec{\omega}) \quad := \quad \frac{1}{\left(\sqrt{2\,\pi}\right)^n} \int_{\mathbb{C}^n} f(\vec{x})\, e^{+\sigma\, i\, \vec{x}\, \vec{\omega}} d^n x \tag{20.2-1}$$

where $\sigma = \pm 1$ is the sign of the transform. The FT is a unitary transform. Its inverse is the complex conjugate transform

$$f(\vec{x}) \quad = \quad \frac{1}{\left(\sqrt{2\,\pi}\right)^n} \int_{\mathbb{C}^n} F(\vec{\omega})\, e^{-\sigma\, i\, \vec{x}\, \vec{\omega}} d^n \omega \tag{20.2-2}$$

For the 1-dimensional case one has

$$F(\omega) \quad = \quad \frac{1}{\sqrt{2\,\pi}} \int_{-\infty}^{+\infty} f(x) e^{+\sigma\, i\, x\, \omega} dx \tag{20.2-3}$$

$$f(x) \quad = \quad \frac{1}{\sqrt{2\,\pi}} \int_{-\infty}^{+\infty} F(\omega)\, e^{-\sigma\, i\, x\, \omega} d\omega \tag{20.2-4}$$

The 'frequency-form' is

$$F(\nu) \quad = \quad \int_{-\infty}^{+\infty} f(x) e^{+2\,\pi\, \sigma\, i\, x\, \nu} dx \tag{20.2-5}$$

$$f(x) \quad = \quad \int_{-\infty}^{+\infty} F(\nu)\, e^{-2\,\pi\, \sigma\, i\, x\, \nu} d\nu \tag{20.2-6}$$

The semi-continuous Fourier transform

For periodic functions defined on a interval $L \in \mathbb{R}$, $f : L \to \mathbb{R}, \; x \mapsto f(x)$ one defines the *semi-continuous Fourier transform* as:

$$c_k \quad := \quad \frac{1}{\sqrt{L}} \int_L f(x)\, e^{+\sigma\, 2\,\pi\, i\, k\, x/L} dx \tag{20.2-7}$$

The inverse transform is an infinite sum with the property

$$\frac{1}{\sqrt{L}} \sum_{k=-\infty}^{k=+\infty} c_k\, e^{-\sigma\, 2\,\pi\, i\, k\, x/L} \quad = \quad \begin{cases} f(x) & \text{if } f \text{ continuous at } x \\ \frac{f(x+0)+f(x-0)}{2} & \text{else} \end{cases} \tag{20.2-8}$$

An equivalent form of the semi-continuous Fourier transform is given by

$$a_k \quad := \quad \frac{1}{\sqrt{L}} \int_L f(x)\, \cos \frac{2\,\pi\, k\, x}{L} dx, \qquad k = 0, 1, 2, \ldots \tag{20.2-9a}$$

$$b_k \quad := \quad \frac{1}{\sqrt{L}} \int_L f(x)\, \sin \frac{2\,\pi\, k\, x}{L} dx, \qquad k = 1, 2, \ldots \tag{20.2-9b}$$

$$f(x) \quad = \quad \frac{1}{\sqrt{L}} \left[\frac{a_0}{2} + \sum_{k=1}^{\infty} \left(a_k\, \cos \frac{2\,\pi\, k\, x}{L} + b_k\, \sin \frac{2\,\pi\, k\, x}{L} \right) \right] \tag{20.2-9c}$$

The connection between the two forms is given by:

$$c_k \quad = \quad \begin{cases} \frac{a_0}{2} & (k = 0) \\ \frac{1}{2}(a_k - i b_k) & (k > 0) \\ \frac{1}{2}(a_k + i b_k) & (k < 0) \end{cases} \tag{20.2-10}$$

The discrete Fourier transform

The *discrete Fourier transform* (DFT) of a sequence f of length n with elements f_x is defined by

$$c_k := \frac{1}{\sqrt{n}} \sum_{x=0}^{n-1} f_x \, e^{+\sigma \, 2 \, \pi \, i \, x \, k / n} \tag{20.2-11}$$

The inverse transform is

$$f_x = \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} c_k \, e^{-\sigma \, 2 \, \pi \, i \, x \, k / n} \tag{20.2-12}$$

Some sources define $c_k := \sum \ldots$ and $f_x = \frac{1}{n} \sum \ldots$. With this definition the transform does not preserve norms but there is a computational saving when the normalization does not matter. When a transform and its inverse are applied in succession (as with convolution algorithms) then only one normalization (factors $\frac{1}{n}$) is needed instead of two (factors $\frac{1}{\sqrt{n}}$). In the implementations of the transform we will in general omit the normalization altogether and leave it to the user to apply it where needed.

20.3 Radix-2 FFT algorithms

A little bit of notation

In what follows let a be a length-n sequence with n a power of two.

- Let $a^{(even)}$ and $a^{(odd)}$ denote the length-$n/2$ subsequences of those elements of a that have even and odd indices, respectively. That is, $a^{(even)} = [a_0, a_2, a_4, a_6, \ldots, a_{n-2}]$ and $a^{(odd)} = [a_1, a_3, \ldots, a_{n-1}]$.
- Let $a^{(left)}$ and $a^{(right)}$ denote the left and right subsequences, respectively. That is, $a^{(left)} = [a_0, a_1, \ldots, a_{n/2-1}]$ and $a^{(right)} = [a_{n/2}, a_{n/2+1}, \ldots, a_{n-1}]$.
- Let $c = \mathcal{S}^k a$ denote the sequence with elements $c_x = a_x \, e^{\sigma \, 2 \, \pi \, i \, k \, x / n}$ where $\sigma = \pm 1$ is the sign of the transform. The symbol $\mathcal{S}$ shall suggest a shift operator. With radix-2 FFT algorithms only $\mathcal{S}^{1/2}$ is needed.
- In relations between sequences we sometimes emphasize the length of the sequences on both sides as in $a^{(even)} \overset{n/2}{=} b^{(odd)} + c^{(odd)}$. In these relations the operators plus and minus are to be understood as element-wise.

20.3.1 Decimation in time (DIT) FFT

The following observation is the key to the (radix-2) *decimation in time* (DIT) FFT algorithm, also called *Cooley-Tukey* FFT algorithm: For even values of n the k-th element of the Fourier transform is

$$\mathcal{F}\big[\, a \,\big]_k = \sum_{x=0}^{n-1} a_x \, z^{x \, k} = \sum_{x=0}^{n/2-1} a_{2\,x} \, z^{2\,x\,k} + \sum_{x=0}^{n/2-1} a_{2\,x+1} \, z^{(2\,x+1)\,k} \tag{20.3-1a}$$

$$= \sum_{x=0}^{n/2-1} a_{2\,x} \, z^{2\,x\,k} + z^k \sum_{x=0}^{n/2-1} a_{2\,x+1} \, z^{2\,x\,k} \tag{20.3-1b}$$

where $z = e^{\sigma \, i \, 2 \, \pi / n}$, $\sigma = \pm 1$ is the sign of the transform, and $k \in \{0, 1, \ldots, n-1\}$.

The identity tells us how to compute the k-th element of the length-n Fourier transform from the length-$n/2$ Fourier transforms of the even and odd indexed subsequences.

To actually rewrite the length-n FT in terms of length-$n/2$ FTs one has to distinguish whether $0 \le k < n/2$ or $n/2 \le k < n$. In the expressions we rewrite $k \in \{0, 1, 2, \ldots, n-1\}$ as $k = j + \delta\,\frac{n}{2}$ where $j \in \{0, 1, 2, \ldots, n/2-1\}$ and $\delta \in \{0, 1\}$:

$$\sum_{x=0}^{n-1} a_x\, z^{x\,(j+\delta\,\frac{n}{2})} = \sum_{x=0}^{n/2-1} a_x^{(even)}\, z^{2\,x\,(j+\delta\,\frac{n}{2})} + z^{j+\delta\,\frac{n}{2}} \sum_{x=0}^{n/2-1} a_x^{(odd)}\, z^{2\,x\,(j+\delta\,\frac{n}{2})} \tag{20.3-2a}$$

$$= \begin{cases} \displaystyle\sum_{x=0}^{n/2-1} a_x^{(even)}\, z^{2\,x\,j} + z^{j} \sum_{x=0}^{n/2-1} a_x^{(odd)}\, z^{2\,x\,j} & \text{for} \quad \delta = 0 \\ \displaystyle\sum_{x=0}^{n/2-1} a_x^{(even)}\, z^{2\,x\,j} - z^{j} \sum_{x=0}^{n/2-1} a_x^{(odd)}\, z^{2\,x\,j} & \text{for} \quad \delta = 1 \end{cases} \tag{20.3-2b}$$

The minus sign in the relation for $\delta = 1$ is due to the fact that $z^{j+1\cdot n/2} = z^j\, z^{n/2} = -z^j$.

Observing that z^2 is just the root of unity that appears in a length-$n/2$ transform one can rewrite the last two equations to obtain the *radix-2 DIT FFT step*:

$$\mathcal{F}\big[\,a\,\big]^{(left)} \stackrel{n/2}{=} \mathcal{F}\big[\,a^{(even)}\,\big] + \mathcal{S}^{1/2}\mathcal{F}\big[\,a^{(odd)}\,\big] \tag{20.3-3a}$$

$$\mathcal{F}\big[\,a\,\big]^{(right)} \stackrel{n/2}{=} \mathcal{F}\big[\,a^{(even)}\,\big] - \mathcal{S}^{1/2}\mathcal{F}\big[\,a^{(odd)}\,\big] \tag{20.3-3b}$$

The length-n transform has been replaced by two transforms of length $n/2$. If n is a power of 2 this scheme can be applied recursively until length-one transforms are reached which are identity ('do nothing') operations.

Thereby the operation count is improved to proportional $n \cdot \log_2(n)$: there are $\log_2(n)$ splitting steps, the work in each step is proportional to n.

Note that the operator $\mathcal{S}$ depends on the sign of the transform.

20.3.1.1 Recursive implementation

A recursive implementation of radix-2 DIT FFT given as pseudo code (C++ version in [FXT: fft/recfft2.cc]) is

```
1   procedure rec_fft_dit2(a[], n, x[], is)
2   // complex a[0..n-1] input
3   // complex x[0..n-1] result
4   {
5       complex b[0..n/2-1], c[0..n/2-1]   // workspace
6       complex s[0..n/2-1], t[0..n/2-1]   // workspace
7
8       if n == 1 then  // end of recursion
9       {
10          x[0] := a[0]
11          return
12      }
13
14      nh := n/2
15
16      for k:=0 to nh-1  // copy to workspace
17      {
18          s[k] := a[2*k]    // even indexed elements
19          t[k] := a[2*k+1]  // odd  indexed elements
20      }
21
22      // recursion: call two half-length FFTs:
23      rec_fft_dit2(s[], nh, b[], is)
24      rec_fft_dit2(t[], nh, c[], is)
25
26      fourier_shift(c[], nh, is*1/2)
27
```

```
28      for k:=0 to nh-1  // copy back from workspace
29      {
30          x[k]    := b[k] + c[k]
31          x[k+nh] := b[k] - c[k]
32      }
33  }
```

The parameter `is` $= \sigma = \pm 1$ is the sign of the transform. The data length `n` must be a power of 2. The result is returned in the array `x[]`. Note that normalization (multiplication of each element of `x[]` by $1/\sqrt{n}$) is not included here.

The procedure uses the subroutine `fourier_shift` which modifies the array `c[]` according to the operation $\mathcal{S}^v$: each element `c[k]` is multiplied by $e^{v\,2\,\pi\,i\,k/n}$. It is called with $v = \pm 1/2$ for the Fourier transform. The pseudo code (C++ equivalent in [FXT: fft/fouriershift.cc]) is

```
1   procedure fourier_shift(c[], n, v)
2   {
3       for k:=0 to n-1
4       {
5           c[k] := c[k] * exp(v*2.0*PI*I*k/n)
6       }
7   }
```

The recursive FFT-procedure involves $n \log_2(n)$ function calls, which can be avoided by rewriting it in a non-recursive way. One can even do all operations *in-place*, no temporary workspace is needed at all. The price is the necessity of an additional data reordering: the procedure `revbin_permute(a[],n)` rearranges the array `a[]` in a way that each element a_x is swapped with $a_{\tilde{x}}$, where $\tilde{x}$ is obtained from x by reversing its binary digits. Methods to do this are discussed in section 2.1.

20.3.1.2 Iterative implementation

Pseudo code for a non-recursive procedure of the radix-2 DIT algorithm (C++ implementation given in [FXT: fft/fftdit2.cc]):

```
1   procedure fft_depth_first_dit2(a[], ldn, is)
2   // complex a[0..2**ldn-1] input, result
3   {
4       n := 2**ldn  // length of a[] is a power of 2
5
6       revbin_permute(a[], n)
7
8       for ldm:=1 to ldn  // log_2(n) iterations
9       {
10          m  := 2**ldm
11          mh := m/2
12
13          for r:=0 to n-m step m  // n/m iterations
14          {
15              for j:=0 to mh-1  // m/2 iterations
16              {
17                  e := exp(is*2*PI*I*j/m)  // log_2(n)*n/m*m/2 = log_2(n)*n/2 computations
18
19                  u := a[r+j]
20                  v := a[r+j+mh] * e
21
22                  a[r+j]    := u + v
23                  a[r+j+mh] := u - v
24              }
25          }
26      }
27  }
```

This version of a non-recursive FFT procedure already avoids the calling overhead and it works in-place. But it is a bit wasteful. The (expensive) computation `e := exp(is*2*PI*I*j/m)` is done $n/2 \cdot \log_2(n)$ times.

20.3.1.3 Saving trigonometric computations

To reduce the number of sine and cosine computations, one can simply swap the two inner loops, leading to the first 'real world' FFT procedure presented here. Pseudo code for a non-recursive procedure of the radix-2 DIT algorithm (C++ version in [FXT: fft/fftdit2.cc]):

```
1   procedure fft_dit2(a[], ldn, is)
2   // complex a[0..2**ldn-1] input, result
3   {
4       n := 2**ldn
5
6       revbin_permute(a[], n)
7
8       for ldm:=1 to ldn  // log_2(n) iterations
9       {
10          m  := 2**ldm
11          mh := m/2
12
13          for j:=0 to mh-1  // m/2 iterations
14          {
15              e := exp(is*2*PI*I*j/m)  // 1 + 2 + ... + n/8 + n/4 + n/2 == n-1 computations
16
17              for r:=0 to n-m step m
18              {
19                  u := a[r+j]
20                  v := a[r+j+mh] * e
21
22                  a[r+j]    := u + v
23                  a[r+j+mh] := u - v
24              }
25          }
26      }
27  }
```

Swapping the two inner loops reduces the number of trigonometric computations to `n` but leads to a feature that many FFT implementations share: memory access is highly nonlocal. For each recursion stage (value of `ldm`) the array is traversed `mh` times with `n/m` accesses in strides[1] of `mh`. As `mh` is a power of two this can (on computers that use memory cache) have a very negative performance impact for large values of `n`. On computers where the memory access is very slow compared to the CPU the naive version can actually be faster.

It is a good idea to extract the `ldm==1` stage of the outermost loop, this avoids complex multiplications with the trivial factors $1 + 0\,i$ (and the computations of these quantities as trigonometric functions). Replace

```
for ldm:=1 to ldn
{
```

by

```
for r:=0 to n-1 step 2
{
    {a[r], a[r+1]} := {a[r]+a[r+1], a[r]-a[r+1]}
}
for ldm:=2 to ldn
{
```

20.3.2 Decimation in frequency (DIF) FFT

Splitting of the Fourier sum into a left and right half leads to the *decimation in frequency* (DIF) FFT algorithm, also called *Sande-Tukey* FFT algorithm. For even values of n the k-th element of the Fourier

[1] Stride-n memory access: consecutive accesses to addresses that are n units apart.

transform is

$$\mathcal{F}\big[\,a\,\big]_k \;=\; \sum_{x=0}^{n-1} a_x\, z^{x\,k} \;=\; \sum_{x=0}^{n/2-1} a_x\, z^{x\,k} + \sum_{x=n/2}^{n-1} a_x\, z^{x\,k} \tag{20.3-4a}$$

$$= \sum_{x=0}^{n/2-1} a_x\, z^{x\,k} + \sum_{x=0}^{n/2-1} a_{x+n/2}\, z^{(x+n/2)\,k} \tag{20.3-4b}$$

$$= \sum_{x=0}^{n/2-1} (a_x^{(left)} + z^{k\,n/2}\, a_x^{(right)})\, z^{x\,k} \tag{20.3-4c}$$

where $z = e^{\sigma\, i\, 2\,\pi/n}$, $\sigma = \pm 1$ is the sign of the transform, and $k \in \{0, 1, \ldots, n-1\}$.

Here one has to distinguish whether k is even or odd. Therefore we rewrite $k \in \{0, 1, 2, \ldots, n-1\}$ as $k = 2\,j + \delta$ where $j \in \{0, 1, 2, \ldots, n/2-1\}$ and $\delta \in \{0, 1\}$:

$$\sum_{x=0}^{n-1} a_x\, z^{x\,(2\,j+\delta)} \;=\; \sum_{x=0}^{n/2-1} (a_x^{(left)} + z^{(2\,j+\delta)\,n/2}\, a_x^{(right)})\, z^{x\,(2\,j+\delta)} \tag{20.3-5a}$$

$$= \begin{cases} \displaystyle\sum_{x=0}^{n/2-1} (a_x^{(left)} + a_x^{(right)})\, z^{2\,x\,j} & \text{for} \quad \delta = 0 \\ \displaystyle\sum_{x=0}^{n/2-1} z^x (a_x^{(left)} - a_x^{(right)})\, z^{2\,x\,j} & \text{for} \quad \delta = 1 \end{cases} \tag{20.3-5b}$$

$z^{(2\,j+\delta)\,n/2} = e^{\pm\pi\, i\, \delta}$ is equal to plus or minus one for $\delta = 0$ or $\delta = 1$ corresponding to k even or odd. The last two equations are, more compactly written, the key to the *radix-2 DIF FFT step*:

$$\mathcal{F}\big[\,a\,\big]^{(even)} \;\overset{n/2}{=}\; \mathcal{F}\big[\,a^{(left)} + a^{(right)}\,\big] \tag{20.3-6a}$$

$$\mathcal{F}\big[\,a\,\big]^{(odd)} \;\overset{n/2}{=}\; \mathcal{F}\big[\,\mathcal{S}^{1/2}\left(a^{(left)} - a^{(right)}\right)\big] \tag{20.3-6b}$$

A recursive implementation of radix-2 DIF FFT given as pseudo code (C++ version given in [FXT: fft/recfft2.cc]) is

```
1   procedure rec_fft_dif2(a[], n, x[], is)
2   // complex a[0..n-1] input
3   // complex x[0..n-1] result
4   {
5       complex b[0..n/2-1], c[0..n/2-1]   // workspace
6       complex s[0..n/2-1], t[0..n/2-1]   // workspace
7
8       if n == 1 then
9       {
10          x[0] := a[0]
11          return
12      }
13
14      nh := n/2
15
16      for k:=0 to nh-1
17      {
18          s[k] := a[k]     // 'left'  elements
19          t[k] := a[k+nh]  // 'right' elements
20      }
21
22      for k:=0 to nh-1
23      {
24          {s[k], t[k]} := {(s[k]+t[k]), (s[k]-t[k])}
25      }
26
27      fourier_shift(t[], nh, is*0.5)
28
29      rec_fft_dif2(s[], nh, b[], is)
30      rec_fft_dif2(t[], nh, c[], is)
```

```
31
32      j := 0
33      for k:=0 to nh-1
34      {
35          x[j]   := b[k]
36          x[j+1] := c[k]
37          j := j+2
38      }
39  }
```

The parameter `is` $= \sigma = \pm 1$ is the sign of the transform. The data length `n` must be a power of 2. The result is returned in the array `x[]`. Again, the routine does no normalization.

Pseudo code for a non-recursive procedure (a C++ implementation is given in [FXT: fft/fftdif2.cc]):

```
1   procedure fft_dif2(a[],ldn,is)
2   // complex a[0..2**ldn-1] input, result
3   {
4       n := 2**ldn
5
6       for ldm:=ldn to 1 step -1
7       {
8           m  := 2**ldm
9           mh := m/2
10
11          for j:=0 to mh-1
12          {
13              e := exp(is*2*PI*I*j/m)
14
15              for r:=0 to n-m step m
16              {
17                  u := a[r+j]
18                  v := a[r+j+mh]
19
20                  a[r+j]    := (u + v)
21                  a[r+j+mh] := (u - v) * e
22              }
23          }
24      }
25
26      revbin_permute(a[], n)
27  }
```

In DIF FFTs the `revbin_permute()`-procedure is called after the main loop, in the DIT code it is called before the main loop. As in the procedure for the DIT FFT (section 20.3.1.3 on page 399) the inner loops where swapped to save trigonometric computations.

Extracting the `ldm==1` stage of the outermost loop is again a good idea. Replace the line

```
for  ldm:=ldn to 1 step -1
```

by

```
for  ldm:=ldn to 2 step -1
```

and insert

```
for r:=0 to n-1 step 2
{
    {a[r], a[r+1]} := {a[r]+a[r+1], a[r]-a[r+1]}
}
```

before the call of `revbin_permute(a[], n)`.

20.4 Saving trigonometric computations

The sine and cosine computations are an expensive part of any FFT. There are two apparent ways for saving the involved CPU cycles, the use of lookup-tables and recursive methods. The CORDIC algorithms for sine and cosine given in section 33.2.1 on page 650 can be useful when implementing FFTs in hardware.

20.4.1 Using lookup tables

The idea is to precompute all necessary values, store them in an array, and later looking up the values needed. This is a good idea if one wants to compute many FFTs of the same (small) length. For FFTs of large sequences one gets large lookup tables that can introduce a high cache-miss rate. Thereby one is likely experiencing little or no speed gain, even a notable slowdown is possible. However, for a length-n FFT one does not need to store all the (n complex or $2\,n$ real) sine/cosine values $\exp(2\,\pi\,i\,k/n) = \cos(2\,\pi\,k/n) + i\,\sin(2\,\pi\,k/n)$ where $k = 0, 1, 2, 3, \ldots, n-1$. For the lookups one can use the symmetry relations

$$\cos(\pi + x) = -\cos(x) \tag{20.4-1a}$$
$$\sin(\pi + x) = -\sin(x) \tag{20.4-1b}$$

To reduce the interval from $0 \ldots 2\pi$ to $0 \ldots \pi$. Exploiting the relations

$$\cos(\pi/2 + x) = -\sin(x) \tag{20.4-2a}$$
$$\sin(\pi/2 + x) = +\cos(x) \tag{20.4-2b}$$

further reduces the interval to $0 \ldots \pi/2$. Finally, the relation

$$\sin(x) = \cos(\pi/2 - x) \tag{20.4-3}$$

shows that only the table of cosines is needed. That is, already a table of the $n/4$ real values $\cos(2\,\pi\,i\,k/n)$ for $k = 0, 1, 2, 3, \ldots, n/4-1$ suffices for a length-n FFT computation. The size of the table is thereby cut by a factor of 8. Possible cache problems can sometimes be mitigated by simply storing the trigonometric values in reversed order which can avoid many equidistant memory accesses.

20.4.2 Recursive generation

In FFT computations one typically needs the values

$$[\exp(i\,\varphi\,0) = 1, \quad \exp(i\,\varphi\,\gamma), \quad \exp(i\,\varphi\,2\,\gamma), \quad \exp(i\,\varphi\,3\,\gamma), \quad \ldots] \quad \text{where} \quad \varphi \in \mathbb{R}$$

in sequence. The naive idea for a recursive computation of these values is to precompute $d = \exp(i\,\varphi\,\gamma)$ and then compute the following value using the identity $\exp(i\,\varphi\,k\,\gamma) = d \cdot \exp(i\,\varphi\,(k-1)\,\gamma)$. This method, however, is of no practical value because the numerical error grows exponentially in the process.

A stable version of a trigonometric recursion for the computation of the sequence can be stated as follows. Precompute

$$c = \cos\varphi, \tag{20.4-4a}$$
$$s = \sin\varphi, \tag{20.4-4b}$$
$$\alpha = 1 - \cos\gamma \qquad \text{[Cancellation!]} \tag{20.4-4c}$$
$$= 2\left(\sin\frac{\gamma}{2}\right)^2 \qquad \text{[OK.]} \tag{20.4-4d}$$
$$\beta = \sin\gamma \tag{20.4-4e}$$

Then compute the next pair (c_+, s_+) from the previous one (c, s) via

$$c_+ = c - (\alpha\,c + \beta\,s); \tag{20.4-5a}$$
$$s_+ = s - (\alpha\,s - \beta\,c); \tag{20.4-5b}$$

The underlying idea is to use the relation $E(\varphi + \gamma) = E(\varphi) - E(\varphi) \cdot z$ where $E(x) := \exp(2\,\pi\,i\,x)$. This leads to $z = 1 - \cos\gamma - i\,\sin\gamma = 2\,\left(\sin\frac{\gamma}{2}\right)^2 - i\,\sin\gamma$.

Do not expect to get all the precision you would get with calls of the sine and cosine functions, but even for very long FFTs less than 3 bits of precision are lost. When working with (C-type) `double`s it might be a good idea to use the type `long double` with the trigonometric recursion: the generated values will then always be accurate within the `double`-precision, provided `long double`s are actually more precise than `double`s. With high precision multiplication routines (that is, with exact integer convolution) this can be mandatory.

A real-world example from [FXT: `fht_dif_core()` in fht/fhtdif.cc]:

```
1           [--snip--]
2           double tt = M_PI_4/kh;   // the angle increment
3           double s1 = 0.0,  c1 = 1.0;  // start at angle zero
4           double al = sin(0.5*tt);
5           al *= (2.0*al);
6           double be = sin(tt);
7
8           for (ulong i=1; i<kh; i++)
9           {
10              double t1 = c1;
11              c1 -= (al*t1+be*s1);
12              s1 -= (al*s1-be*t1);
13
14              // here c1 = cos(tt*i) and s1 = sin(tt*i)
15          [--snip--]
```

The variable `tt` equals γ in relations 20.4-4d and 20.4-4e.

20.5 Higher radix FFT algorithms

Higher radix FFT algorithms save trigonometric computations. The radix-4 FFT algorithms presented in what follows replace all multiplications with complex factors $(0, \pm i)$ by the obvious simpler operations. Radix-8 algorithms also simplify the special cases where the sines and cosines equal $\pm\sqrt{1/2}$.

Further the bookkeeping overhead is reduced due to the more unrolled structure. Moreover, the number of loads and stores is reduced.

More notation

Let a be a length-n sequence where n is a multiple of m.

- Let $a^{(r\%m)}$ denote the subsequence of the elements with index x where $x \equiv r \bmod m$. For example, $a^{(0\%2)} = a^{(even)}$ and $a^{(3\%4)} = [a_3, a_7, a_{11}, a_{15}, \dots]$. The length of $a^{(r\%m)}$ is n/m.
- Let $a^{(r/m)}$ denote the subsequence of elements with indices $\frac{r\,n}{m} \dots \frac{(r+1)\,n}{m} - 1$. For example $a^{(1/2)} = a^{(right)}$ and $a^{(2/3)}$ is the last third of a. The length of $a^{(r/m)}$ is also n/m.

20.5.1 Decimation in time algorithms

First rewrite the radix-2 DIT step (relations 20.3-3a and 20.3-3b on page 397) in the new notation:

$$\mathcal{F}\big[\,a\,\big]^{(0/2)} \overset{n/2}{=} \mathcal{S}^{0/2}\mathcal{F}\big[\,a^{(0\%2)}\,\big] + \mathcal{S}^{1/2}\mathcal{F}\big[\,a^{(1\%2)}\,\big] \tag{20.5-1a}$$

$$\mathcal{F}\big[\,a\,\big]^{(1/2)} \overset{n/2}{=} \mathcal{S}^{0/2}\mathcal{F}\big[\,a^{(0\%2)}\,\big] - \mathcal{S}^{1/2}\mathcal{F}\big[\,a^{(1\%2)}\,\big] \tag{20.5-1b}$$

The operator $\mathcal{S}$ is defined on section 20.3 on page 396, note that $\mathcal{S}^{0/2} = \mathcal{S}^0$ is the identity operator.

The derivation of the radix-4 step is analogous to the radix-2 step, it just involves more writing and does not give additional insights. So we just state the *radix-4 DIT FFT step* which can be applied when n is

divisible by 4:

$$\mathcal{F}\big[\,a\,\big]^{(0/4)} \overset{n/4}{=} +\mathcal{S}^{0/4}\mathcal{F}\big[\,a^{(0\%4)}\,\big] + \mathcal{S}^{1/4}\mathcal{F}\big[\,a^{(1\%4)}\,\big] + \mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(2\%4)}\,\big] + \mathcal{S}^{3/4}\mathcal{F}\big[\,a^{(3\%4)}\,\big] \quad (20.5\text{-}2a)$$

$$\mathcal{F}\big[\,a\,\big]^{(1/4)} \overset{n/4}{=} +\mathcal{S}^{0/4}\mathcal{F}\big[\,a^{(0\%4)}\,\big] + i\sigma\mathcal{S}^{1/4}\mathcal{F}\big[\,a^{(1\%4)}\,\big] - \mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(2\%4)}\,\big] - i\sigma\mathcal{S}^{3/4}\mathcal{F}\big[\,a^{(3\%4)}\,\big] \quad (20.5\text{-}2b)$$

$$\mathcal{F}\big[\,a\,\big]^{(2/4)} \overset{n/4}{=} +\mathcal{S}^{0/4}\mathcal{F}\big[\,a^{(0\%4)}\,\big] - \mathcal{S}^{1/4}\mathcal{F}\big[\,a^{(1\%4)}\,\big] + \mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(2\%4)}\,\big] - \mathcal{S}^{3/4}\mathcal{F}\big[\,a^{(3\%4)}\,\big] \quad (20.5\text{-}2c)$$

$$\mathcal{F}\big[\,a\,\big]^{(3/4)} \overset{n/4}{=} +\mathcal{S}^{0/4}\mathcal{F}\big[\,a^{(0\%4)}\,\big] - i\sigma\mathcal{S}^{1/4}\mathcal{F}\big[\,a^{(1\%4)}\,\big] - \mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(2\%4)}\,\big] + i\sigma\mathcal{S}^{3/4}\mathcal{F}\big[\,a^{(3\%4)}\,\big] \quad (20.5\text{-}2d)$$

In contrast to the radix-2 step that happens to be identical for forward and backward transform the sign of the transform $\sigma = \pm 1$ appears explicitly. The relations, written more compactly:

$$\begin{aligned}\mathcal{F}\big[\,a\,\big]^{(j/4)} \overset{n/4}{=}\ & +e^{\sigma\,2\,i\,\pi\,0\,j/4}\cdot\mathcal{S}^{0/4}\mathcal{F}\big[\,a^{(0\%4)}\,\big] + e^{\sigma\,2\,i\,\pi\,1\,j/4}\cdot\mathcal{S}^{1/4}\mathcal{F}\big[\,a^{(1\%4)}\,\big] \\ & +e^{\sigma\,2\,i\,\pi\,2\,j/4}\cdot\mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(2\%4)}\,\big] + e^{\sigma\,2\,i\,\pi\,3\,j/4}\cdot\mathcal{S}^{3/4}\mathcal{F}\big[\,a^{(3\%4)}\,\big]\end{aligned} \quad (20.5\text{-}3)$$

where $j \in \{0,1,2,3\}$ and n is a multiple of 4. Still more compactly:

$$\mathcal{F}\big[\,a\,\big]^{(j/4)} \overset{n/4}{=} \sum_{k=0}^{3} e^{\sigma 2\,i\,\pi\,k\,j/4}\cdot\mathcal{S}^{k/4}\mathcal{F}\big[\,a^{(k\%4)}\,\big] \qquad j \in \{0,1,2,3\} \quad (20.5\text{-}4)$$

where the summation symbol denotes *element-wise* summation of the sequences. The dot indicates multiplication of all elements of the sequence by the exponential.

The general *radix-r DIT FFT step*, applicable when n is a multiple of r, is:

$$\mathcal{F}\big[\,a\,\big]^{(j/r)} \overset{n/r}{=} \sum_{k=0}^{r-1} e^{\sigma\,2\,i\,\pi\,k\,j/r}\cdot\mathcal{S}^{k/r}\mathcal{F}\big[\,a^{(k\%r)}\,\big] \qquad j = 0,1,2,\ldots,r-1 \quad (20.5\text{-}5)$$

Our notation turned out to be useful indeed.

20.5.2 Decimation in frequency algorithms

The radix-2 DIF step (relations 20.3-6a and 20.3-6b on page 400), in the new notation:

$$\mathcal{F}\big[\,a\,\big]^{(0\%2)} \overset{n/2}{=} \mathcal{F}\big[\,\mathcal{S}^{0/2}\Big(a^{(0/2)} + a^{(1/2)}\Big)\,\big] \quad (20.5\text{-}6a)$$

$$\mathcal{F}\big[\,a\,\big]^{(1\%2)} \overset{n/2}{=} \mathcal{F}\big[\,\mathcal{S}^{1/2}\Big(a^{(0/2)} - a^{(1/2)}\Big)\,\big] \quad (20.5\text{-}6b)$$

The *radix-4 DIF FFT step*, applicable for n divisible by 4, is

$$\mathcal{F}\big[\,a\,\big]^{(0\%4)} \overset{n/4}{=} \mathcal{F}\big[\,\mathcal{S}^{0/4}\Big(a^{(0/4)} + a^{(1/4)} + a^{(2/4)} + a^{(3/4)}\Big)\,\big] \quad (20.5\text{-}7a)$$

$$\mathcal{F}\big[\,a\,\big]^{(1\%4)} \overset{n/4}{=} \mathcal{F}\big[\,\mathcal{S}^{1/4}\Big(a^{(0/4)} + i\,\sigma\,a^{(1/4)} - a^{(2/4)} - i\,\sigma\,a^{(3/4)}\Big)\,\big] \quad (20.5\text{-}7b)$$

$$\mathcal{F}\big[\,a\,\big]^{(2\%4)} \overset{n/4}{=} \mathcal{F}\big[\,\mathcal{S}^{2/4}\Big(a^{(0/4)} - a^{(1/4)} + a^{(2/4)} - a^{(3/4)}\Big)\,\big] \quad (20.5\text{-}7c)$$

$$\mathcal{F}\big[\,a\,\big]^{(3\%4)} \overset{n/4}{=} \mathcal{F}\big[\,\mathcal{S}^{3/4}\Big(a^{(0/4)} - i\,\sigma\,a^{(1/4)} - a^{(2/4)} + i\,\sigma\,a^{(3/4)}\Big)\,\big] \quad (20.5\text{-}7d)$$

Again, $\sigma = \pm 1$ is the sign of the transform. More compactly:

$$\mathcal{F}\big[\,a\,\big]^{(j\%4)} \overset{n/4}{=} \mathcal{F}\big[\,\mathcal{S}^{j/4}\sum_{k=0}^{3} e^{\sigma\,2\,i\,\pi\,k\,j/4}\cdot a^{(k/4)}\,\big] \qquad j \in \{0,1,2,3\} \quad (20.5\text{-}8)$$

The general *radix-r DIF FFT step* is

$$\mathcal{F}\big[\,a\,\big]^{(j\%r)} \overset{n/r}{=} \mathcal{F}\big[\,\mathcal{S}^{j/r}\sum_{k=0}^{r-1} e^{\sigma\,2\,i\,\pi\,k\,j/r}\cdot a^{(k/r)}\,\big] \qquad j \in \{0,1,2,\ldots,r-1\} \quad (20.5\text{-}9)$$

20.5.3 Implementation of radix-r FFTs

For the implementation of a radix-r FFT with $r \neq 2$ the `revbin_permute` routine has to be replaced by its radix-r version `radix_permute`. The reordering now swaps elements a_x with $a_{\tilde{x}}$ where $\tilde{x}$ is obtained from x by reversing its radix-r expansion (see section 2.2 on page 97). In most practical cases one considers $r = p^x$ where p is prime. Pseudo code for a radix $r = p^x$ DIT FFT:

```
1   procedure fftdit_r(a[], n, is)
2   // complex a[0..n-1] input, result
3   // p (hardcoded)
4   // r == power of p (hardcoded)
5   // n == power of p (not necessarily a power of r)
6   {
7       radix_permute(a[], n, p)
8
9       lx := log(r) / log(p)  // r == p ** lx
10      ln := log(n) / log(p)
11      ldm := (log(n)/log(p)) % lx
12
13      if ( ldm != 0 )  // n is not a power of p
14      {
15          xx := p**lx
16          for z:=0 to n-xx step xx
17          {
18              fft_dit_xx(a[z..z+xx-1], is)  // inlined length-xx dit fft
19          }
20      }
21
22      for ldm:=ldm+lx to ln step lx
23      {
24          m  := p**ldm
25          mr := m/r
26
27          for j := 0 to mr-1
28          {
29              e := exp(is*2*PI*I*j/m)
30
31              for k:=0 to n-m step m
32              {
33                  // all code in this block should be
34                  // inlined, unrolled and combined:
35
36                  // temporary  u[0..r-1]
37
38                  for z:=0 to r-1
39                  {
40                      u[z] := a[k+j+mr*z]
41                  }
42
43                  radix_permute(u[], r, p)
44
45                  for z:=1 to r-1  // e**0 == 1
46                  {
47                      u[z] := u[z] * e**z
48                  }
49
50                  r_point_fft(u[], is)
51
52                  for z:=0 to r-1
53                  {
54                      a[k+j+mr*z] := u[z]
55                  }
56              }
57          }
58      }
59  }
```

Of course the loops that use the variable `z` have to be unrolled, the (length-p^x) scratch space `u[]` has to be replaced by explicit variables (for example, `u0, u1, ...`), and the `r_point_fft(u[],is)` shall be an inlined p^x-point FFT.

There is one pitfall: if one uses the radix-p permutation instead of a radix-p^x permutation (for example, the radix-2 `revbin_permute()` for a radix-4 FFT), then some additional reordering is necessary in the innermost loop. In the given pseudo code this is indicated by the `radix_permute(u[],p)` just before the

`p_point_fft(u[],is)` line.

20.5.4 Radix-4 DIT FFT

C++ code for a radix-4 DIT FFT is given in [FXT: fft/fftdit4l.cc]:

```
static const ulong RX = 4;  // == r
static const ulong LX = 2;  // == log(r)/log(p) == log_2(r)

void
fft_dit4l(Complex *f, ulong ldn, int is)
// Decimation in time radix-4 FFT.
{
    double s2pi = ( is>0 ? 2.0*M_PI : -2.0*M_PI );

    const ulong n = (1UL<<ldn);

    revbin_permute(f, n);

    ulong ldm = (ldn&1);

    if ( ldm!=0 )  // n is not a power of 4, need a radix-2 step
    {
        for (ulong r=0; r<n; r+=2)
        {
            Complex a0 = f[r];
            Complex a1 = f[r+1];

            f[r]   = a0 + a1;
            f[r+1] = a0 - a1;
        }
    }

    ldm += LX;

    for ( ; ldm<=ldn ; ldm+=LX)
    {
        ulong m = (1UL<<ldm);
        ulong m4 = (m>>LX);
        double ph0 = s2pi/m;

        for (ulong j=0; j<m4; j++)
        {
            double phi = j*ph0;
            Complex e  = SinCos(phi);
            Complex e2 = SinCos(2.0*phi);
            Complex e3 = SinCos(3.0*phi);

            for (ulong r=0; r<n; r+=m)
            {
                ulong i0 = j + r;
                ulong i1 = i0 + m4;
                ulong i2 = i1 + m4;
                ulong i3 = i2 + m4;

                Complex a0 = f[i0];
                Complex a1 = f[i2]; // (!)
                Complex a2 = f[i1]; // (!)
                Complex a3 = f[i3];

                a1 *= e;
                a2 *= e2;
                a3 *= e3;

                Complex t0 = (a0+a2) + (a1+a3);
                Complex t2 = (a0+a2) - (a1+a3);

                Complex t1 = (a0-a2) + Complex(0,is) * (a1-a3);
                Complex t3 = (a0-a2) - Complex(0,is) * (a1-a3);

                f[i0] = t0;
                f[i1] = t1;
                f[i2] = t2;
                f[i3] = t3;
            }
```

```
70          }
71      }
72  }
```

For reasonable performance the call to the procedure `radix_permute(u[],p)` of the pseudo code has been replaced by changing indices in the loops where the `a[z]` are read. The respective lines are marked with the comment '`// (!)`'.

In order not to restrict the possible array sizes to powers of $p^x = 4$ but only to powers of $p = 2$ an additional radix-2 step has been prepended that is used when n is an odd power of two.

The routine [FXT: `fft_dit4_core_p1()` in fft/fftdit4.cc] is a reasonably optimized radix-4 DIT FFT implementation. It starts with an radix-2 or radix-8 step for the initial pass with trivial exp()-values. The core routine is hardcoded for $\sigma = +1$ and called with swapped real and imaginary part for the inverse transform as explained in section 20.8 on page 415.

The routine, however, uses separate arrays for real and imaginary parts which is very problematic with large transforms: the memory access pattern in skips that are a power of two *will* lead to cache problems. A routine that uses the C++ type `complex` is given in [FXT: fft/cfftdit4.cc]. The core routine is hardcoded for $\sigma = -1$ (therefore the name suffix `_m1`):

```
1   void
2   fft_dit4_core_m1(Complex *f, ulong ldn)
3   // Auxiliary routine for fft_dit4()
4   // Radix-4 decimation in frequency FFT
5   // ldn := base-2 logarithm of the array length
6   // Fixed isign = -1
7   // Input data must be in revbin_permuted order
8   {
9       const ulong n = (1UL<<ldn);
10
11      if ( n<=2 )
12      {
13          if ( n==2 )  sumdiff(f[0], f[1]);
14          return;
15      }
16
17      ulong ldm = ldn & 1;
18      if ( ldm!=0 )  // n is not a power of 4, need a radix-8 step
19      {
20          for (ulong i0=0; i0<n; i0+=8)  fft8_dit_core_m1(f+i0);  // isign
21      }
22      else
23      {
24          for (ulong i0=0; i0<n; i0+=4)
25          {
26              ulong i1 = i0 + 1;
27              ulong i2 = i1 + 1;
28              ulong i3 = i2 + 1;
29
30              Complex x, y, u, v;
31              sumdiff(f[i0], f[i1], x, u);
32              sumdiff(f[i2], f[i3], y, v);
33              v *= Complex(0, -1);  // isign
34              sumdiff(u, v, f[i1], f[i3]);
35              sumdiff(x, y, f[i0], f[i2]);
36          }
37      }
38      ldm += 2 * LX;
39
40
41      for ( ; ldm<=ldn; ldm+=LX)
42      {
43          ulong m = (1UL<<ldm);
44          ulong m4 = (m>>LX);
45          const double ph0 = -2.0*M_PI/m;  // isign
46
47          for (ulong j=0; j<m4; j++)
48          {
49              double phi = j * ph0;
50              Complex e  = SinCos(phi);
51              Complex e2 = e * e;
52              Complex e3 = e2 * e;
```

```
53
54              for (ulong r=0; r<n; r+=m)
55              {
56                  ulong i0 = j + r;
57                  ulong i1 = i0 + m4;
58                  ulong i2 = i1 + m4;
59                  ulong i3 = i2 + m4;
60
61                  Complex x = f[i1] * e2;
62                  Complex u;
63                  sumdiff3_r(x, f[i0], u);
64
65                  Complex v = f[i3] * e3;
66                  Complex y = f[i2] * e;
67                  sumdiff(y, v);
68                  v *= Complex(0, -1);  // isign
69
70                  sumdiff(u, v, f[i1], f[i3]);
71                  sumdiff(x, y, f[i0], f[i2]);
72              }
73          }
74      }
75  }
```

The `sumdiff()` function is defined in [FXT: aux0/sumdiff.h]:

```
1   template <typename Type>
2   static inline void sumdiff(Type &a, Type &b)
3   // {a, b}  <--| {a+b, a-b}
4   { Type t=a-b; a+=b; b=t; }
```

The routine `fft8_dit_core_m1()` is an unrolled size-8 DIT FFT (hardcoded for $\sigma = -1$) given in [FXT: fft/fft8ditcore.cc]. We further need a version of the routine for the positive sign. It uses a routine `fft8_dit_core_p1()` for the computation of length-8 DIT FFTs (with $\sigma = -1$). The following changes need to be made in the core routine [FXT: fft/cfftdit4.cc]:

```
1   void
2   fft_dit4_core_p1(Complex *f, ulong ldn)
3   // Fixed isign = +1
4   {
5       [--snip--]
6           for (ulong i0=0; i0<n; i0+=8)  fft8_dit_core_p1(f+i0);  // isign
7       [--snip--]
8               v *= Complex(0, +1);  // isign
9       [--snip--]
10          const double ph0 = +2.0*M_PI/m;  // isign
11      [--snip--]
12                  v *= Complex(0, +1);  // isign
13      [--snip--]
14  }
```

The routine to be called by the user is

```
1   void
2   fft_dit4(Complex *f, ulong ldn, int is)
3   // Fast Fourier Transform
4   // ldn := base-2 logarithm of the array length
5   // is := sign of the transform (+1 or -1)
6   // Radix-4 decimation in time algorithm
7   {
8       revbin_permute(f, 1UL<<ld);
9       if ( is>0 )  fft_dit4_core_p1(f, ldn);
10      else         fft_dit4_core_m1(f, ldn);
11  }
```

A version that uses the separate arrays for real and imaginary part is given in [FXT: fft/fftdit4.cc]. The type `complex` version, should be preferred for large transforms.

20.5.5 Radix-4 DIF FFT

Pseudo code for a radix-4 DIF FFT:

```
procedure fftdif4(a[], ldn, is)
// complex a[0..2**ldn-1] input, result
{
    n := 2**ldn

    for ldm := ldn to 2 step -2
    {
        m  := 2**ldm
        mr := m/4

        for j := 0 to mr-1
        {
            e  := exp(is*2*PI*I*j/m)
            e2 := e * e
            e3 := e2 * e

            for r := 0 to n-m step m
            {
                u0 := a[r+j]
                u1 := a[r+j+mr]
                u2 := a[r+j+mr*2]
                u3 := a[r+j+mr*3]

                x := u0 + u2
                y := u1 + u3
                t0 := x + y  // == (u0+u2) + (u1+u3)
                t2 := x - y  // == (u0+u2) - (u1+u3)

                x := u0 - u2
                y := (u1 - u3)*I*is
                t1 := x + y  // == (u0-u2) + (u1-u3)*I*is
                t3 := x - y  // == (u0-u2) - (u1-u3)*I*is

                t1 := t1 * e
                t2 := t2 * e2
                t3 := t3 * e3

                a[r+j]      := t0
                a[r+j+mr]   := t2  // (!)
                a[r+j+mr*2] := t1  // (!)
                a[r+j+mr*3] := t3
            }
        }
    }

    if is_odd(ldn) then  // n not a power of 4
    {
        for r:=0 to n-2 step 2
        {
            {a[r], a[r+1]} := {a[r]+a[r+1], a[r]-a[r+1]}
        }
    }

    revbin_permute(a[],n)
}
```

The C++ equivalent is [FXT: `fft_dif4l()` in fft/fftdif4l.cc]. A reasonably optimized implementation is given in [FXT: fft/cfftdif4.cc], it is hardcoded for $\sigma = +1$ (therefore the name suffix `_p1`):

```
static const ulong RX = 4;
static const ulong LX = 2;

void
fft_dif4_core_p1(Complex *f, ulong ldn)
// Auxiliary routine for fft_dif4().
// Radix-4 decimation in frequency FFT.
// Output data is in revbin_permuted order.
// ldn := base-2 logarithm of the array length.
// Fixed isign = +1
{
    const ulong n = (1UL<<ldn);

    if ( n<=2 )
    {
        if ( n==2 )  sumdiff(f[0], f[1]);
        return;
    }

    for (ulong ldm=ldn; ldm>=(LX<<1); ldm-=LX)
    {
```

```
22          ulong m = (1UL<<ldm);
23          ulong m4 = (m>>LX);
24
25          const double ph0 = 2.0*M_PI/m;  // isign
26
27          for (ulong j=0; j<m4; j++)
28          {
29              double phi = j * ph0;
30              Complex e  = SinCos(phi);
31              Complex e2 = e * e;
32              Complex e3 = e2 * e;
33
34              for (ulong r=0; r<n; r+=m)
35              {
36                  ulong i0 = j + r;
37                  ulong i1 = i0 + m4;
38                  ulong i2 = i1 + m4;
39                  ulong i3 = i2 + m4;
40
41                  Complex x, y, u, v;
42                  sumdiff(f[i0], f[i2], x, u);
43                  sumdiff(f[i1], f[i3], y, v);
44                  v *= Complex(0, +1);  // isign
45
46                  diffsum3(x, y, f[i0]);
47                  f[i1] = y * e2;
48
49                  sumdiff(u, v, x, y);
50                  f[i3] = y * e3;
51                  f[i2] = x * e;
52              }
53          }
54      }
55
56
57      if ( ldn & 1 )  // n is not a power of 4, need a radix-8 step
58      {
59          for (ulong i0=0; i0<n; i0+=8)  fft8_dif_core_p1(f+i0);  // isign
60      }
61      else
62      {
63          for (ulong i0=0; i0<n; i0+=4)
64          {
65              ulong i1 = i0 + 1;
66              ulong i2 = i1 + 1;
67              ulong i3 = i2 + 1;
68
69              Complex x, y, u, v;
70              sumdiff(f[i0], f[i2], x, u);
71              sumdiff(f[i1], f[i3], y, v);
72              v *= Complex(0, +1);  // isign
73              sumdiff(x, y, f[i0], f[i1]);
74              sumdiff(u, v, f[i2], f[i3]);
75          }
76      }
77  }
```

The routine for $\sigma = -1$ needs changes where the comment `isign` appears [FXT: fft/cfftdif4.cc]::

```
1   void
2   fft_dif4_core_m1(Complex *f, ulong ldn)
3   // Fixed isign = -1
4   {
5       [--snip--]
6           const double ph0 = -2.0*M_PI/m;  // isign
7       [--snip--]
8                   v *= Complex(0, -1);  // isign
9       [--snip--]
10          for (ulong i0=0; i0<n; i0+=8)  fft8_dif_core_m1(f+i0);  // isign
11      [--snip--]
12              v *= Complex(0, -1);  // isign
13      [--snip--]
14  }
```

The routine to be called by the user is

```
1   void
2   fft_dif4(Complex *f, ulong ldn, int is)
3   // Fast Fourier Transform
4   // ldn := base-2 logarithm of the array length
5   // is := sign of the transform (+1 or -1)
6   // radix-4 decimation in frequency algorithm
7   {
8       if ( is>0 )  fft_dif4_core_p1(f, ldn);
9       else         fft_dif4_core_m1(f, ldn);
10      revbin_permute(f, 1UL<<ldn);
11  }
```

A version that uses the separate arrays for real and imaginary part is given in [FXT: fft/fftdif4.cc]. Again, the type `complex` version, should be preferred for large transforms.

20.6 Split-radix Fourier transforms

The idea underlying the *split-radix FFT algorithm* is to use both radix-2 and radix-4 decompositions at the same time. We use one relation from the radix-2 (DIF) decomposition (relation 20.3-6a on page 400, the one for the even indices), and for the odd indices we use the radix-4 splitting (relations 20.5-7b and 20.5-7d on page 404) in a slightly reordered form. The radix-4 decimation in frequency (DIF) step for the split-radix FFT is

$$\mathcal{F}\big[\,a\,\big]^{(0\%2)} \stackrel{n/2}{=} \mathcal{F}\big[\,\Big(a^{(0/2)} + a^{(1/2)}\Big)\,\big] \tag{20.6-1a}$$

$$\mathcal{F}\big[\,a\,\big]^{(1\%4)} \stackrel{n/4}{=} \mathcal{F}\big[\,\mathcal{S}^{1/4}\Big(\Big(a^{(0/4)} - a^{(2/4)}\Big) + i\,\sigma\,\Big(a^{(1/4)} - a^{(3/4)}\Big)\Big)\,\big] \tag{20.6-1b}$$

$$\mathcal{F}\big[\,a\,\big]^{(3\%4)} \stackrel{n/4}{=} \mathcal{F}\big[\,\mathcal{S}^{3/4}\Big(\Big(a^{(0/4)} - a^{(2/4)}\Big) - i\,\sigma\,\Big(a^{(1/4)} - a^{(3/4)}\Big)\Big)\,\big] \tag{20.6-1c}$$

Now we have expressed the length-$N = 2^n$ FFT as one length-$N/2$ and two length-$N/4$ FFTs. A nice feature is that the operation count of the split-radix FFT is actually lower than that of the radix-4 FFT. Using the introduced notation it is almost trivial to write down the DIT version of the algorithm: The radix-4 decimation in time (DIT) step for the split-radix FFT is

$$\mathcal{F}\big[\,a\,\big]^{(0/2)} \stackrel{n/2}{=} \Big(\mathcal{F}\big[\,a^{(0\%2)}\,\big] + \mathcal{S}^{1/2}\mathcal{F}\big[\,a^{(1\%2)}\,\big]\Big) \tag{20.6-2a}$$

$$\mathcal{F}\big[\,a\,\big]^{(1/4)} \stackrel{n/4}{=} \Big(\mathcal{F}\big[\,a^{(0\%4)}\,\big] - \mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(2\%4)}\,\big]\Big) + i\sigma\mathcal{S}^{1/4}\Big(\mathcal{F}\big[\,a^{(1\%4)}\,\big] - \mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(3\%4)}\,\big]\Big) \tag{20.6-2b}$$

$$\mathcal{F}\big[\,a\,\big]^{(3/4)} \stackrel{n/4}{=} \Big(\mathcal{F}\big[\,a^{(0\%4)}\,\big] - \mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(2\%4)}\,\big]\Big) - i\sigma\mathcal{S}^{1/4}\Big(\mathcal{F}\big[\,a^{(1\%4)}\,\big] - \mathcal{S}^{2/4}\mathcal{F}\big[\,a^{(3\%4)}\,\big]\Big) \tag{20.6-2c}$$

Pseudo code for the split-radix DIF algorithm:

```
1   procedure fft_splitradix_dif(x[], y[], ldn, is)
2   {
3       n := 2**ldn
4
5       if n<=1  return
6
7       n2 := 2*n
8
9       for k:=1 to ldn
10      {
11          n2 := n2 / 2
12          n4 := n2 / 4
13
14          e := 2 * PI / n2
15
16          for j:=0 to n4-1
17          {
18              a := j * e
19              cc1 := cos(a)
20              ss1 := sin(a)
21              cc3 := cos(3*a)  // == 4*cc1*(cc1*cc1-0.75)
22              ss3 := sin(3*a)  // == 4*ss1*(0.75-ss1*ss1)
```

```
23
24              ix := j
25              id := 2*n2
26
27              while ix<n-1
28              {
29                  i0 := ix
30                  while i0 < n
31                  {
32                      i1 := i0 + n4
33                      i2 := i1 + n4
34                      i3 := i2 + n4
35
36                      {x[i0], r1} := {x[i0] + x[i2], x[i0] - x[i2]}
37                      {x[i1], r2} := {x[i1] + x[i3], x[i1] - x[i3]}
38
39                      {y[i0], s1} := {y[i0] + y[i2], y[i0] - y[i2]}
40                      {y[i1], s2} := {y[i1] + y[i3], y[i1] - y[i3]}
41
42                      {r1, s3} := {r1+s2, r1-s2}
43                      {r2, s2} := {r2+s1, r2-s1}
44
45                      // complex mult: (x[i2],y[i2]) := -(s2,r1) * (ss1,cc1)
46                      x[i2] :=  r1*cc1 - s2*ss1
47                      y[i2] := -s2*cc1 - r1*ss1
48
49                      // complex mult: (y[i3],x[i3]) := (r2,s3) * (cc3,ss3)
50                      x[i3] :=  s3*cc3 + r2*ss3
51                      y[i3] :=  r2*cc3 - s3*ss3
52
53                      i0 := i0 + id
54                  }
55
56                  ix := 2 * id - n2 + j
57                  id := 4 * id
58              }
59          }
60      }
61
62      ix := 1
63      id := 4
64
65      while ix<n
66      {
67          for i0:=ix-1 to n-id step id
68          {
69              i1 := i0 + 1
70              {x[i0], x[i1]} := {x[i0]+x[i1], x[i0]-x[i1]}
71              {y[i0], y[i1]} := {y[i0]+y[i1], y[i0]-y[i1]}
72          }
73
74          ix := 2 * id - 1
75          id := 4 * id
76      }
77
78      revbin_permute(x[],n)
79      revbin_permute(y[],n)
80
81      if is>0
82      {
83          for j:=1 to n/2-1
84          {
85              swap(x[j], x[n-j])
86          }
87
88          for j:=1 to n/2-1
89          {
90              swap(y[j], y[n-j])
91          }
92      }
93  }
```

The C++ implementation given in [FXT: fft/fftsplitradix.cc] uses a DIF core as above which was given in [106]. The C++ type `complex` version of the split-radix FFT given in [FXT: fft/cfftsplitradix.cc] uses a DIF or DIT core, depending on the sign of the transform. Here we just give the DIF version:

```
1   void
2   split_radix_dif_fft_core(Complex *f, ulong ldn)
3   // Split-radix decimation in frequency (DIF) FFT.
4   // ldn := base-2 logarithm of the array length.
5   // Fixed isign = +1
6   // Output data is in revbin_permuted order.
```

```
 7  {
 8      if ( ldn==0 )  return;
 9
10      const ulong n = (1UL<<ldn);
11
12      double s2pi = 2.0*M_PI;  // pi*2*isign
13      ulong n2 = 2*n;
14      for (ulong k=1; k<ldn; k++)
15      {
16          n2 >>= 1;  // == n>>(k-1) == n, n/2, n/4, ..., 4
17          const ulong n4 = n2 >> 2;  // == n/4, n/8, ..., 1
18          const double e = s2pi / n2;
19
20          {  // j==0:
21              const ulong j = 0;
22              ulong ix = j;
23              ulong id = (n2<<1);
24              while ( ix<n )
25              {
26                  for (ulong i0=ix; i0<n; i0+=id)
27                  {
28                      ulong i1 = i0 + n4;
29                      ulong i2 = i1 + n4;
30                      ulong i3 = i2 + n4;
31
32                      Complex t0, t1;
33                      sumdiff3(f[i0], f[i2], t0);
34                      sumdiff3(f[i1], f[i3], t1);
35
36                      // t1 *= Complex(0, 1);  // +isign
37                      t1 = Complex(-t1.imag(), t1.real());
38
39                      sumdiff(t0, t1);
40                      f[i2] = t0; // * Complex(cc1, ss1);
41                      f[i3] = t1; // * Complex(cc3, ss3);
42                  }
43
44                  ix = (id<<1) - n2 + j;
45                  id <<= 2;
46              }
47          }
48
49          for (ulong j=1; j<n4; j++)
50          {
51              double a = j * e;
52              double cc1,ss1, cc3,ss3;
53              SinCos(a, &ss1, &cc1);
54              SinCos(3.0*a, &ss3, &cc3);
55
56              ulong ix = j;
57              ulong id = (n2<<1);
58              while ( ix<n )
59              {
60                  for (ulong i0=ix; i0<n; i0+=id)
61                  {
62                      ulong i1 = i0 + n4;
63                      ulong i2 = i1 + n4;
64                      ulong i3 = i2 + n4;
65
66                      Complex t0, t1;
67                      sumdiff3(f[i0], f[i2], t0);
68                      sumdiff3(f[i1], f[i3], t1);
69
70                      t1 = Complex(-t1.imag(), t1.real());
71
72                      sumdiff(t0, t1);
73                      f[i2] = t0 * Complex(cc1, ss1);
74                      f[i3] = t1 * Complex(cc3, ss3);
75                  }
76
77                  ix = (id<<1) - n2 + j;
78                  id <<= 2;
79              }
80          }
81      }
82
83      for (ulong ix=0, id=4;  ix<n;  id*=4)
```

```
84      {
85          for (ulong i0=ix; i0<n; i0+=id)  sumdiff(f[i0], f[i0+1]);
86          ix = 2*(id-1);
87      }
88  }
```

The function `sumdiff3()` is defined in [FXT: aux0/sumdiff.h]:

```
1   template <typename Type>
2   static inline void sumdiff3(Type &a, Type b, Type &d)
3   // {a, b, d} <--| {a+b, b, a-b}  (used in split-radix fft)
4   { d=a-b; a+=b; }
```

20.7 Symmetries of the Fourier transform

A bit of notation again. Let $\overline{a}$ be the length-n sequence a reversed around the element with index 0:

$$\begin{aligned} \overline{a}_0 &:= a_0 && (20.7\text{-}1a)\\ \overline{a}_{n/2} &:= a_{n/2} \quad \text{if } n \text{ even} && (20.7\text{-}1b)\\ \overline{a}_k &:= a_{n-k} = a_{-k} && (20.7\text{-}1c) \end{aligned}$$

That is, we consider the indices modulo n and $\overline{a}$ is the sequence a with negated indices. Element zero stays in its place and for even n there is also an element with index $n/2$ that stays in place.
Example one, length-4: $a := [0, 1, 2, 3]$ then $\overline{a} = [0, 3, 2, 1]$ (zero and two stay).
Example two, length-5: $a := [0, 1, 2, 3, 4]$ then $\overline{a} = [0, 4, 3, 2, 1]$ (only zero stays).

Let a_S and a_A denote the symmetric and antisymmetric part of the sequence a, respectively:

$$\begin{aligned} a_S &:= \frac{1}{2}\,(a + \overline{a}) && (20.7\text{-}2a)\\ a_A &:= \frac{1}{2}\,(a - \overline{a}) && (20.7\text{-}2b) \end{aligned}$$

The elements with index 0 (and $n/2$ for even n) a_A are zero. One has

$$\begin{aligned} a &= a_S + a_A && (20.7\text{-}3a)\\ \overline{a} &= a_S - a_A && (20.7\text{-}3b) \end{aligned}$$

Let $c + i\,d$ be the FT of the sequence $a + i\,b$. Then

$$\begin{aligned} \mathcal{F}\big[\,(a_S + a_A) + i\,(b_S + b_A)\,\big] &= (c_S + c_A) + i\,(d_S + d_A) \quad \text{where} && (20.7\text{-}4a)\\ \mathcal{F}\big[\,a_S\,\big] &= c_S \quad \in \mathbb{R} && (20.7\text{-}4b)\\ \mathcal{F}\big[\,a_A\,\big] &= i\,d_A \quad \in i\,\mathbb{R} && (20.7\text{-}4c)\\ \mathcal{F}\big[\,i\,b_S\,\big] &= i\,d_S \quad \in i\,\mathbb{R} && (20.7\text{-}4d)\\ \mathcal{F}\big[\,i\,b_A\,\big] &= c_A \quad \in \mathbb{R} && (20.7\text{-}4e) \end{aligned}$$

Where we write $a \in \mathbb{R}$ as a short form for a purely real sequence a. Equivalently, write $a \in i\,\mathbb{R}$ for purely imaginary sequences. Thereby the FT of a complex symmetric or antisymmetric sequence is symmetric or antisymmetric, respectively:

$$\begin{aligned} \mathcal{F}\big[\,a_S + i\,b_S\,\big] &= c_S + i\,d_S && (20.7\text{-}5a)\\ \mathcal{F}\big[\,a_A + i\,b_A\,\big] &= c_A + i\,d_A && (20.7\text{-}5b) \end{aligned}$$

The real and imaginary part of the transform of a symmetric sequence correspond to the real and imaginary part of the original sequence. With an antisymmetric sequence the transform of the real

and imaginary part correspond to the imaginary and real part of the original sequence.

$$\mathcal{F}\big[\,(a_S + a_A)\,\big] \;=\; c_S + i\,d_A \tag{20.7-6a}$$
$$\mathcal{F}\big[\,i\,(b_S + b_A)\,\big] \;=\; c_A + i\,d_S \tag{20.7-6b}$$

Now let the sequence a be purely real. Then

$$\mathcal{F}\big[\,a_S\,\big] \;=\; +\overline{\mathcal{F}\big[\,a_S\,\big]} \qquad \in\ \mathbb{R} \tag{20.7-7a}$$
$$\mathcal{F}\big[\,a_A\,\big] \;=\; -\overline{\mathcal{F}\big[\,a_A\,\big]} \qquad \in i\,\mathbb{R} \tag{20.7-7b}$$

That is, the FT of a real symmetric sequence is real and symmetric and the FT of a real antisymmetric sequence is purely imaginary and antisymmetric. Thereby the FT of a general real sequence is the complex conjugate of its reversed:

$$\mathcal{F}\big[\,a\,\big] \;=\; \overline{\mathcal{F}\big[\,a\,\big]}^{\,*} \qquad for \quad a \in \mathbb{R} \tag{20.7-8}$$

Similarly, for a purely imaginary sequence $b \in i\mathbb{R}$:

$$\mathcal{F}\big[\,b_S\,\big] \;=\; +\overline{\mathcal{F}\big[\,b_S\,\big]} \qquad \in i\,\mathbb{R} \tag{20.7-9a}$$
$$\mathcal{F}\big[\,b_A\,\big] \;=\; -\overline{\mathcal{F}\big[\,b_A\,\big]} \qquad \in\ \mathbb{R} \tag{20.7-9b}$$

We compare the results of the Fourier transform and its inverse (the transform with negated sign σ) by symbolically writing the transforms as a complex multiplication with the trigonometric term (using C for cosine, S for sine):

$$\mathcal{F}\big[\,a + i\,b\,\big] \;:\; (a + i\,b)\,(C + i\,S) \;=\; (a\,C - b\,S) + i\,(b\,C + a\,S) \tag{20.7-10a}$$
$$\mathcal{F}^{-1}\big[\,a + i\,b\,\big] \;:\; (a + i\,b)\,(C - i\,S) \;=\; (a\,C + b\,S) + i\,(b\,C - a\,S) \tag{20.7-10b}$$

The terms on the right side can be identified with those in relation 20.7-4a. We see that changing the sign of the transform leads to a result where the components due to the antisymmetric components of the input are negated.

Now write $\mathcal{F}$ for the Fourier operator, and $\mathcal{R}$ for the reversal. We have $\mathcal{F}^4 = \mathrm{id}$, $\mathcal{F}^3 = \mathcal{F}^{-1}$, and $\mathcal{F}^2 = \mathcal{R}$. Thereby the inverse transform can be computed as either

$$\mathcal{F}^{-1} \;=\; \mathcal{R}\,\mathcal{F} \;=\; \mathcal{F}\,\mathcal{R} \tag{20.7-11}$$

20.8 Inverse FFT for free

Some FFT implementations are hardcoded for a fixed sign of the transform. If one cannot easily modify the implementation into the transform with the other sign (the inverse transform), then how can one compute the inverse FFT?

If the implementation uses separate arrays for real and imaginary part of the complex sequences to be transformed, as in

```
1  procedure my_fft(ar[], ai[], ldn)  // only for is==+1 !
2  // real ar[0..2**ldn-1] input, result, real part
3  // real ai[0..2**ldn-1] input, result, imaginary part
4  {
5      // incredibly complicated code
6      // that you cannot see how to modify
7      // for is==-1
8  }
```

then do a follows: with the forward transform being

```
my_fft(ar[], ai[], ldn)  // forward FFT
```

compute the inverse transform as

```
my_fft(ai[], ar[], ldn)  // backward FFT
```

Note the swapped real and imaginary parts! The same trick works for a procedure coded for fixed `is`$= -1$.

To see why this works, we first note that

$$\begin{aligned} \mathcal{F}\big[\,a + i\,b\,\big] &= \mathcal{F}\big[\,a_S\,\big] + i\,\sigma\,\mathcal{F}\big[\,a_A\,\big] + i\,\mathcal{F}\big[\,b_S\,\big] + \sigma\,\mathcal{F}\big[\,b_A\,\big] && \text{(20.8-1a)} \\ &= \mathcal{F}\big[\,a_S\,\big] + i\,\mathcal{F}\big[\,b_S\,\big] + i\,\sigma\,\big(\mathcal{F}\big[\,a_A\,\big] - i\,\mathcal{F}\big[\,b_A\,\big]\big) && \text{(20.8-1b)} \end{aligned}$$

and the computation with swapped real- and imaginary parts gives

$$\mathcal{F}\big[\,b + i\,a\,\big] = \mathcal{F}\big[\,b_S\,\big] + i\,\mathcal{F}\big[\,a_S\,\big] + i\,\sigma\,\big(\mathcal{F}\big[\,b_A\,\big] - i\,\mathcal{F}\big[\,a_A\,\big]\big) \qquad \text{(20.8-2a)}$$

... but the real and imaginary parts are implicitly swapped at the end of the computation, giving

$$\mathcal{F}\big[\,a_S\,\big] + i\,\mathcal{F}\big[\,b_S\,\big] - i\,\sigma\,\big(\mathcal{F}\big[\,a_A\,\big] - i\,\mathcal{F}\big[\,b_A\,\big]\big) = \mathcal{F}^{-1}\big[\,a + i\,b\,\big] \qquad \text{(20.8-2b)}$$

When a complex type is used then the best way to achieve the inverse transform may be to reverse the sequence according to the symmetry of the FT according to relation 20.7-11: the transform with negated sign can be computed by reversing the order of the result (use [FXT: `reverse_0()` in perm/reverse.h]). The reversal can also happen with the input data before the transform, which is advantageous if the data has to be copied anyway (use [FXT: `copy_reverse_0()` in aux1/copy.h]). While not really 'free' the additional work will usually not matter.

A mechanical way to obtain a routine for the inverse FFT from a given FFT routine for length n is to replace all reads and writes at nonzero array indices i by the operations at indices $n - i$ (where n is identified with zero).

20.9 Real valued Fourier transforms

The Fourier transform of a purely real sequence $c = \mathcal{F}\big[\,a\,\big]$ where $a \in \mathbb{R}$ has a symmetric real part ($\mathfrak{Re}\,\overline{c} = \mathfrak{Re}\,c$, relation 20.7-8) and an antisymmetric imaginary part ($\mathfrak{Im}\,\overline{c} = -\,\mathfrak{Im}\,c$). The symmetric and antisymmetric part of the original sequence correspond to the symmetric (and purely real) and antisymmetric (and purely imaginary) part of the transform, respectively:

$$\mathcal{F}\big[\,a\,\big] = \mathcal{F}\big[\,a_S\,\big] + i\,\sigma\,\mathcal{F}\big[\,a_A\,\big] \qquad \text{(20.9-1)}$$

Simply using a complex FFT for real input is basically a waste by a factor 2 of memory and CPU cycles. There are several ways out:

- wrapper routines for complex FFTs (section 20.9.3 on page 418)
- usage of the fast Hartley transform (section 24.5 on page 509)
- special versions of the split-radix algorithm (section 20.9.4 on page 419)

All techniques have in common that they store only half of the complex result to avoid the redundancy due to the symmetries of a complex FT of purely real input. The result of a real to (half-) complex FT (abbreviated R2CFT) contains the purely real components c_0 (the 'DC-part' of the input signal) and, in case n is even, $c_{n/2}$ (the Nyquist frequency part). The inverse procedure, the (half-) complex to real transform (abbreviated C2RFT) must be compatible to the ordering of the R2CFT.

20.9.1 Sign of the transforms

The sign of the transform can be chosen arbitrarily to be either $+1$ or -1. Note that transform with the 'other sign' is *not* the inverse transform. The R2CFT and its inverse C2RFT must use the same sign.

Some R2CFT and C2RFT implementations are hardcoded for a fixed sign. In order to obtain the R2CFT with the other sign the trick (in the spirit of section 20.8) is to negate the imaginary part after the transform. This of course is not much of a trick at all. In case one has to copy the data anyway before the transform one can exploit the relation

$$\mathcal{F}\big[\,\overline{a}\,\big] \;=\; \mathcal{F}\big[\,a_S\,\big] - i\,\sigma\,\mathcal{F}\big[\,a_A\,\big] \tag{20.9-2}$$

That is, copy the real data in reversed order to get the transform with the other sign. This technique does not involve an extra pass and should be virtually for free.

For the complex-to-real FTs (C2RFT) one has to negate the imaginary part before the transform in order to obtain the inverse transform for the other sign.

20.9.2 Data ordering

Let c be the Fourier transform of the purely real sequence, it is stored in the array `a[]`. The procedures presented in what follows use one of the following schemes for storing the transformed sequence.

A scheme that interleaves real and imaginary parts ('complex ordering') is

$$\begin{aligned}
\mathtt{a[0]} &= \mathfrak{Re}\, c_0 \\
\mathtt{a[1]} &= \mathfrak{Re}\, c_{n/2} \\
\mathtt{a[2]} &= \mathfrak{Re}\, c_1 \\
\mathtt{a[3]} &= \mathfrak{Im}\, c_1 \\
\mathtt{a[4]} &= \mathfrak{Re}\, c_2 \\
\mathtt{a[5]} &= \mathfrak{Im}\, c_2 \\
&\vdots \\
\mathtt{a[n-2]} &= \mathfrak{Re}\, c_{n/2-1} \\
\mathtt{a[n-1]} &= \mathfrak{Im}\, c_{n/2-1}
\end{aligned} \tag{20.9-3}$$

Note the absence of the elements $\mathfrak{Im}\, c_0$ and $\mathfrak{Im}\, c_{n/2}$ which are always zero.

Some routines store the real parts in the lower half, and imaginary parts in upper half. The data in the lower half will always be ordered as follows:

$$\begin{aligned}
\mathtt{a[0]} &= \mathfrak{Re}\, c_0 \\
\mathtt{a[1]} &= \mathfrak{Re}\, c_1 \\
\mathtt{a[2]} &= \mathfrak{Re}\, c_2 \\
&\vdots \\
\mathtt{a[n/2]} &= \mathfrak{Re}\, c_{n/2}
\end{aligned} \tag{20.9-4}$$

For the imaginary part of the result there are two schemes:

Scheme 1 ('parallel ordering') is

$$\begin{aligned} \mathtt{a}[\mathtt{n}/2+1] &= \mathfrak{Im}\, c_1 \\ \mathtt{a}[\mathtt{n}/2+2] &= \mathfrak{Im}\, c_2 \\ \mathtt{a}[\mathtt{n}/2+3] &= \mathfrak{Im}\, c_3 \\ &\vdots \\ \mathtt{a}[\mathtt{n}-1] &= \mathfrak{Im}\, c_{n/2-1} \end{aligned} \tag{20.9-5}$$

Scheme 2 ('antiparallel ordering') is

$$\begin{aligned} \mathtt{a}[\mathtt{n}/2+1] &= \mathfrak{Im}\, c_{n/2-1} \\ \mathtt{a}[\mathtt{n}/2+2] &= \mathfrak{Im}\, c_{n/2-2} \\ \mathtt{a}[\mathtt{n}/2+3] &= \mathfrak{Im}\, c_{n/2-3} \\ &\vdots \\ \mathtt{a}[\mathtt{n}-1] &= \mathfrak{Im}\, c_1 \end{aligned} \tag{20.9-6}$$

20.9.3 Real valued FT via wrapper routines

A simple way to use a complex length-$n/2$ FFT for a real length-n FFT (n even) is to use some post- and preprocessing routines. For a real sequence a one feeds the (half length) complex sequence $f = a^{(even)} + i\, a^{(odd)}$ into a complex FFT. Some post-processing is necessary. This is not the most elegant real FFT available, but it is directly usable to turn complex FFTs of any (even) length into a real-valued FFT.

A C++ implementation of the real to complex FFT (R2CFT) is given in [FXT: realfft/realfftwrap.cc], the sign of the transform is hardcoded to $\sigma = +1$:

```
1   void
2   wrap_real_complex_fft(double *f, ulong ldn)
3   // Real to complex FFT (R2CFT)
4   {
5       if ( ldn==0 )  return;
6
7       fht_fft((Complex *)f, ldn-1, +1);  // cast
8
9       const ulong n = 1UL<<ldn;
10      const ulong nh = n/2, n4 = n/4;
11      const double phi0 = M_PI / nh;
12      for(ulong i=1; i<n4; i++)
13      {
14          ulong i1 = 2 * i;   // re low [2, 4, ..., n/2-2]
15          ulong i2 = i1 + 1;  // im low [3, 5, ..., n/2-1]
16
17          ulong i3 = n - i1;  // re hi  [n-2, n-4, ..., n/2+2]
18          ulong i4 = i3 + 1;  // im hi  [n-1, n-3, ..., n/2+3]
19
20          double f1r, f2i;
21          sumdiff05(f[i3], f[i1], f1r, f2i);
22
23          double f2r, f1i;
24          sumdiff05(f[i2], f[i4], f2r, f1i);
25
26          double c, s;
27          double phi = i*phi0;
28          SinCos(phi, &s, &c);
29
30          double tr, ti;
31          cmult(c, s, f2r, f2i, tr, ti);
32
33          // f[i1] = f1r + tr;  // re low
34          // f[i3] = f1r - tr;  // re hi
35          // =^=
36          sumdiff(f1r, tr, f[i1], f[i3]);
```

```
37
38
39          // f[i4] = is * (ti + f1i);  // im hi
40          // f[i2] = is * (ti - f1i);  // im low
41          // =^=
42          sumdiff( ti,  f1i, f[i4], f[i2]);
43      }
44      sumdiff(f[0], f[1]);
45  }
```

The output is ordered according to relations 20.9-3 on page 417. The same ordering must be used for the input for the inverse routine, the complex to real FFT (C2RFT). Again the sign of the transform is hardcoded to $\sigma = +1$:

```
1   void
2   wrap_complex_real_fft(double *f, ulong ldn)
3   // Complex to real FFT (C2RFT).
4   {
5       if ( ldn==0 )  return;
6
7       const ulong n = 1UL<<ldn;
8       const ulong nh = n/2, n4 = n/4;
9       const double phi0 = -M_PI / nh;
10      for(ulong i=1; i<n4; i++)
11      {
12          ulong i1 = 2 * i;   // re low [2, 4, ..., n/2-2]
13          ulong i2 = i1 + 1;  // im low [3, 5, ..., n/2-1]
14
15          ulong i3 = n - i1;  // re hi  [n-2, n-4, ..., n/2+2]
16          ulong i4 = i3 + 1;  // im hi  [n-1, n-3, ..., n/2+3]
17
18          double f1r, f2i;
19          // double f1r =  f[i1] + f[i3];  // re symm
20          // double f2i =  f[i1] - f[i3];  // re asymm
21          // =^=
22          sumdiff(f[i1], f[i3], f1r, f2i);
23
24          double f2r, f1i;
25          // double f2r = -f[i2] - f[i4];  // im symm
26          // double f1i =  f[i2] - f[i4];  // im asymm
27          // =^=
28          sumdiff(-f[i4], f[i2], f1i, f2r);
29
30          double c, s;
31          double phi = i*phi0;
32          SinCos(phi, &s, &c);
33
34          double tr, ti;
35          cmult(c, s, f2r, f2i, tr, ti);
36
37          // f[i1] = f1r + tr;  // re low
38          // f[i3] = f1r - tr;  // re hi
39          // =^=
40          sumdiff(f1r, tr, f[i1], f[i3]);
41
42          // f[i2] = ti - f1i;  // im low
43          // f[i4] = ti + f1i;  // im hi
44          // =^=
45          sumdiff(ti, f1i, f[i4], f[i2]);
46      }
47      sumdiff(f[0], f[1]);
48
49      if ( nh>=2 )  { f[nh] *= 2.0; f[nh+1] *= 2.0; }
50
51      fht_fft((Complex *)f, ldn-1, -1);  // cast
52  }
```

20.9.4 Real valued split-radix Fourier transforms

We give pseudo code for the split-radix real to complex FFT and its inverse. The C++ implementations are given in [FXT: realfft/realfftsplitradix.cc]. The code given here follows [107], see also [245] (erratum for page 859 of [245]: at the start of the `DO 32` loop replace assignments by `CC1=COS(A)`, `SS1=SIN(A)`, `CC3=COS(A3)`, `SS3=SIN(A3)`).

Recall that the pseudo code

```
{ a, b } := { c, d }
```

is a parallel assignment. For example,

```
{x0, x1} := {x0+x1, x0-x1 }
```

would translate to the C code

```
double s = x0 + x1,  d = x0 - x1;
x0 = s;  x1 = d;
```

Use the function [FXT: `sumdiff()` in aux0/sumdiff.h] for the translation of the parallel assignments.

20.9.5 Real to complex split-radix FFT

Pseudo code for the split-radix R2CFT algorithm, the sign of the transform is hardcoded to $\sigma = -1$.

```
1   procedure r2cft_splitradix_dit(x[], ldn)
2   {
3       n := 2**ldn
4
5       revbin_permute(x[], n);
6
7       ix := 1;
8       id := 4;
9       do
10      {
11          i0 := ix-1
12          while i0<n
13          {
14              i1 := i0 + 1
15              {x[i0], x[i1]} := {x[i0]+x[i1], x[i0]-x[i1]}
16              i0 := i0 + id
17          }
18          ix := 2*id-1
19          id := 4 * id
20      }
21      while ix<n
22
23      n2 := 2
24      nn := n/4
25      while nn!=0
26      {
27          ix := 0
28          n2 := 2*n2
29          id := 2*n2
30          n4 := n2/4
31          n8 := n2/8
32          do  // ix loop
33          {
34              i0 := ix
35              while i0<n
36              {
37                  i1 := i0
38                  i2 := i1 + n4
39                  i3 := i2 + n4
40                  i4 := i3 + n4
41
42                  {t1, x[i4]} := {x[i4]+x[i3], x[i4]-x[i3]}
43                  {x[i1], x[i3]} := {x[i1]+t1, x[i1]-t1}
44
45                  if  n4!=1
46                  {
47                      i1 := i1 + n8
48                      i2 := i2 + n8
49                      i3 := i3 + n8
50                      i4 := i4 + n8
51
52                      t1 := (x[i3]+x[i4]) * sqrt(1/2)
53                      t2 := (x[i3]-x[i4]) * sqrt(1/2)
54
55                      {x[i4], x[i3]} := {x[i2]-t1, -x[i2]-t1}
56                      {x[i1], x[i2]} := {x[i1]+t2,  x[i1]-t2}
57                  }
58
59                  i0 := i0 + id
60              }
61
62              ix := 2*id - n2
63              id := 2*id
64          }
65          while ix<n
66
```

```
67          e := 2.0*PI/n2
68          a := e
69
70          for j:=2 to n8
71          {
72              cc1 := cos(a)
73              ss1 := sin(a)
74              cc3 := cos(3*a)  // == 4*cc1*(cc1*cc1-0.75)
75              ss3 := sin(3*a)  // == 4*ss1*(0.75-ss1*ss1)
76
77              a := j*e
78
79              ix := 0
80              id := 2*n2
81
82              do  // ix-loop
83              {
84                  i0 := ix
85                  while  i0<n
86                  {
87                      i1 := i0 + j - 1
88                      i2 := i1 + n4
89                      i3 := i2 + n4
90                      i4 := i3 + n4
91                      i5 := i0 + n4 - j + 1
92                      i6 := i5 + n4
93                      i7 := i6 + n4
94                      i8 := i7 + n4
95
96                      // complex mult: (t2,t1) := (x[i7],x[i3]) * (cc1,ss1)
97                      t1 := x[i3]*cc1 + x[i7]*ss1
98                      t2 := x[i7]*cc1 - x[i3]*ss1
99
100                     // complex mult: (t4,t3) := (x[i8],x[i4]) * (cc3,ss3)
101                     t3 := x[i4]*cc3 + x[i8]*ss3
102                     t4 := x[i8]*cc3 - x[i4]*ss3
103
104                     t5 := t1 + t3
105                     t6 := t2 + t4
106                     t3 := t1 - t3
107                     t4 := t2 - t4
108
109                     {t2, x[i3]} := {t6+x[i6], t6-x[i6]}
110                     x[i8] := t2
111                     {t2, x[i7]} := {x[i2]-t3, -x[i2]-t3}
112                     x[i4] := t2
113                     {t1, x[i6]} := {x[i1]+t5, x[i1]-t5}
114                     x[i1] := t1
115                     {t1, x[i5]} := {x[i5]+t4, x[i5]-t4}
116                     x[i2] := t1
117
118                     i0 := i0 + id
119                 }
120
121                 ix := 2*id - n2
122                 id := 2*id
123
124             }
125             while ix<n
126         }
127         nn := nn/2
128     }
129 }
```

The ordering of the output is given as relations 20.9-4 on page 417 for the real part, and relation 20.9-6 for the imaginary part.

20.9.6 Complex to real split-radix FFT

The following routine is the inverse of `r2cft_splitradix_dit()`. The imaginary part of the input data must be ordered according to relation 20.9-6 on page 418. We give pseudo code for the split-radix C2RFT algorithm, the sign of the transform is hardcoded to $\sigma = -1$:

```
1   procedure c2rft_splitradix_dif(x[], ldn)
2   {
3       n := 2**ldn
4
5       n2 := n/2
6       nn := n/4
7       while nn!=0
8       {
9           ix := 0
```

```
        id := n2
        n2 := n2/2
        n4 := n2/4
        n8 := n2/8
        do  // ix loop
        {
            i0 := ix
            while i0<n
            {
                i1 := i0
                i2 := i1 + n4
                i3 := i2 + n4
                i4 := i3 + n4

                {x[i1], t1} := {x[i1]+x[i3], x[i1]-x[i3]}
                x[i2] := 2*x[i2]
                x[i4] := 2*x[i4]
                {x[i3], x[i4]} := {t1+x[i4], t1-x[i4]}

                if  n4!=1
                {
                    i1 := i1 + n8
                    i2 := i2 + n8
                    i3 := i3 + n8
                    i4 := i4 + n8

                    {x[i1], t1} := {x[i2]+x[i1], x[i2]-x[i1]}
                    {t2, x[i2]} := {x[i4]+x[i3], x[i4]-x[i3]}
                    x[i3] := -sqrt(2)*(t2+t1)
                    x[i4] :=  sqrt(2)*(t1-t2)
                }

                i0 := i0 + id
            }

            ix := 2*id - n2
            id := 2*id
        }
        while ix<n

        e := 2.0*PI/n2
        a := e

        for j:=2 to n8
        {
            cc1 := cos(a)
            ss1 := sin(a)
            cc3 := cos(3*a)  // == 4*cc1*(cc1*cc1-0.75)
            ss3 := sin(3*a)  // == 4*ss1*(0.75-ss1*ss1)
            a := j*e

            ix := 0
            id := 2*n2
            do  // ix-loop
            {
                i0 := ix
                while  i0<n
                {
                    i1 := i0 + j - 1
                    i2 := i1 + n4
                    i3 := i2 + n4
                    i4 := i3 + n4
                    i5 := i0 + n4 - j + 1
                    i6 := i5 + n4
                    i7 := i6 + n4
                    i8 := i7 + n4

                    {x[i1], t1} := {x[i1]+x[i6], x[i1]-x[i6]}
                    {x[i5], t2} := {x[i5]+x[i2], x[i5]-x[i2]}
                    {t3, x[i6]} := {x[i8]+x[i3], x[i8]-x[i3]}
                    {t4, x[i2]} := {x[i4]+x[i7], x[i4]-x[i7]}
                    {t1, t5} := {t1+t4, t1-t4}
                    {t2, t4} := {t2+t3, t2-t3}

                    // complex mult: (x[i7],x[i3]) := (t5,t4)  * (ss1,cc1)
                    x[i3] := t5*cc1 + t4*ss1
                    x[i7] := -t4*cc1 + t5*ss1

                    // complex mult: (x[i4],x[i8]) := (t1,t2)  * (cc3,ss3)
                    x[i4] := t1*cc3 - t2*ss3
                    x[i8] := t2*cc3 + t1*ss3

                    i0 := i0 + id
                }

                ix := 2*id - n2
                id := 2*id
```

```
97              }
98              while ix<n
99          }
100
101         nn := nn/2
102     }
103
104     ix := 1;
105     id := 4;
106     do
107     {
108         i0 := ix-1
109         while i0<n
110         {
111             i1 := i0 + 1
112             {x[i0], x[i1]} := {x[i0]+x[i1], x[i0]-x[i1]}
113             i0 := i0 + id
114         }
115         ix := 2*id-1
116         id := 4 * id
117     }
118     while ix<n
119
120     revbin_permute(x[], n);
121 }
```

20.10 Multidimensional Fourier transforms

20.10.1 Definition

Let $a_{x,y}$ ($x = 0, 1, 2, \ldots, C-1$ and $y = 0, 1, 2, \ldots, R-1$) be a 2-dimensional array. That is, a $R \times C$ 'matrix' of R rows (of length C) and C columns (of length R). Its 2-dimensional Fourier transform is defined by:

$$c = \mathcal{F}\big[\, a \,\big] \tag{20.10-1a}$$

$$c_{k,h} := \frac{1}{\sqrt{n}} \sum_{x=0}^{C-1} \sum_{y=0}^{R-1} a_{x,y}\, z^{+(x\,k/C + y\,h/R)} \quad \text{where} \quad z = e^{\sigma\, 2\,\pi\, i} \tag{20.10-1b}$$

where $k \in \{0, 1, 2, \ldots, C-1\}$, $h \in \{0, 1, 2, \ldots, R-1\}$, and $n = R \cdot C$. The inverse transform is

$$a = \mathcal{F}^{-1}\big[\, c \,\big] \tag{20.10-2a}$$

$$a_{x,y} = \frac{1}{\sqrt{n}} \sum_{k=0}^{C-1} \sum_{h=0}^{R-1} c_{k,h}\, z^{-(x\,k/C + y\,h/R)} \tag{20.10-2b}$$

For a m-dimensional array $a_{\vec{x}}$ (where $\vec{x} = (x_1, x_2, x_3, \ldots, x_m)$ and $x_i \in 0, 1, 2, \ldots, S_i$) the m-dimensional Fourier transform $c_{\vec{k}}$ (where $\vec{k} = (k_1, k_2, k_3, \ldots, k_m)$ and $k_i \in 0, 1, 2, \ldots, S_i$) is defined as

$$c_{\vec{k}} := \frac{1}{\sqrt{n}} \sum_{x_1=0}^{S_1-1} \sum_{x_2=0}^{S_2-1} \cdots \sum_{x_m=0}^{S_m-1} a_{\vec{x}}\, z^{(x_1\,k_1/S_1 + x_2\,k_2/S_2 + \ldots + x_m\,k_m/S_m)} \tag{20.10-3a}$$

The inverse transform is, like in the 1-dimensional case, the complex conjugate transform.

20.10.2 The row-column algorithm

The equation of the definition of the two dimensional FT (relation 20.10-1a) can be recast as

$$c_{k,h} = \frac{1}{\sqrt{n}} \sum_{y=0}^{R-1} \left[\exp\left(y\,h/R\right) \sum_{x=0}^{C-1} a_{x,y}\, \exp\left(x\,k/C\right) \right] \tag{20.10-4}$$

which shows that the 2-dimensional FT can be obtained by first applying 1-dimensional transforms on the rows and then applying 1-dimensional transforms on the columns. The same result is obtained when the columns are transformed first and then the rows.

This leads us directly to the *row-column algorithm* for 2-dimensional FFTs. Pseudo code to compute the two dimensional FT of `a[][]` using the row-column method:

```
1   procedure rowcol_ft(a[][], R, C, is)
2   {
3       complex a[R][C]  // R (length-C) rows, C (length-R) columns
4
5       for r:=0 to R-1  // FFT rows
6       {
7           fft(a[r][], C, is)
8       }
9
10      complex t[R]     // scratch array for columns
11      for c:=0 to C-1  // FFT columns
12      {
13          copy a[0,1,...,R-1][c] to t[]  // get column
14          fft(t[], R, is)
15          copy t[] to a[0,1,...,R-1][c]  // write back column
16      }
17  }
```

Here it is assumed that the rows lie in contiguous memory (as in the C language). The equivalent C++ code is given in [FXT: fft/twodimfft.cc].

Transposing the array before the column pass avoids the notorious problem with memory-cache and will do good for the performance in most cases. That is:

```
1   procedure rowcol_fft2d(a[][], R, C, is)
2   {
3       complex a[R][C]  // R (length-C) rows, C (length-R) columns
4
5       for r:=0 to R-1  // FFT rows
6       {
7           fft(a[r][], C, is)
8       }
9
10      transpose( a[R][C] )  // in-place
11
12      for c:=0 to C-1  // FFT columns (which are rows now)
13      {
14          fft(a[c][], R, is)
15      }
16
17      transpose( a[C][R] )  // transpose back (note swapped R,C)
18  }
```

The transposing back at the end of the routine can be avoided if a back-transform will follow immediately as typical for convolution. The back-transform must then be called with `R` and `C` swapped.

The generalization to higher dimensions is straightforward, C++ code is given in [FXT: fft/ndimfft.cc].

20.11 The matrix Fourier algorithm (MFA)

The matrix Fourier algorithm (MFA) is an algorithm for 1-dimensional FFTs that works for data lengths $n = R\,C$. It is quite similar to the row-column algorithm (relation 20.10-4) for 2-dimensional FFTs. The only differences are n multiplications with trigonometric factors and a final matrix transposition.

Consider the input array as a $R \times C$-matrix (R rows, C columns), the rows shall be contiguous in memory. Then the *matrix Fourier algorithm* (MFA) can be stated as follows:

1. Apply a (length R) FFT on each column.
2. Multiply each matrix element (index r, c) by $\exp(\sigma\, 2\,\pi\, i\, r\, c/n)$
3. Apply a (length C) FFT on each row.

4. Transpose the matrix.

Note the elegance! A variant of the MFA is called *four step FFT* in [23]. A trivial modification is obtained if the steps are executed in reversed order. The *transposed matrix Fourier algorithm* (TMFA) for the FFT:

1. Transpose the matrix.
2. Apply a (length C) FFT on each row of the matrix.
3. Multiply each matrix element (index r, c) by $\exp(\sigma\, 2\,\pi\, i\, r\, c/n)$.
4. Apply a (length R) FFT on each column of the matrix.

A variant of the MFA that, apart from the transpositions, accesses the memory only in consecutive address ranges can be stated as

1. Transpose the matrix.
2. Apply a (length C) FFT on each row of the transposed matrix.
3. Multiply each matrix element (index r, c) by $\exp(\sigma\, 2\,\pi\, i\, r\, c/n)$.
4. Transpose the matrix back.
5. Apply a (length R) FFT on each row of the matrix.
6. Transpose the matrix (if the order of the transformed data matters).

The 'transposed' version of this algorithm is identical. The performance will depend critically on the performance of the transposition routine.

It is usually a good idea to use factors of the data length n that are close to $\sqrt{n}$. Of course one can apply the same algorithm for the row (or column) FFTs again: it can be an improvement to split n into 3 factors (as close to $n^{1/3}$ as possible) if a length-$n^{1/3}$ FFT fits completely into the cache. Especially for systems where CPU clock speed is much higher than memory clock speed the performance may increase drastically, a speedup by a factor of three (even when compared to otherwise very well optimized FFTs) can sometimes be observed. Another algorithm that is efficient with large arrays is the localized transform as described in section 24.9 on page 515 for the Hartley transform.

Chapter 21

Algorithms for fast convolution

This chapter gives several FFT based algorithms for fast convolution. These are in practice the most important applications of the FFT. An efficient algorithm for the convolution of arrays that do not fit into the main memory (mass storage convolution) is given for both complex and real data. Further, weighted convolutions and their algorithms are introduced.

We describe how fast convolution can be used for computing the z-transform of sequences of arbitrary length. Another convolution based algorithm for the Fourier transform of arrays of prime length, Rader's algorithm, is described at the end of the chapter.

Convolution algorithms based on the fast Hartley transform are described in section 24.7. The dyadic convolution, which is computed via the Walsh transform is treated in section 22.7.

21.1 Convolution

The *cyclic convolution* (or *circular convolution*) of two length-n sequences $a = [a_0, a_1, \ldots, a_{n-1}]$ and $b = [b_0, b1, \ldots, b_{n-1}]$ is defined as the length-n sequence h with elements h_τ as:

$$h = a \circledast b \tag{21.1-1a}$$

$$h_\tau := \sum_{x+y \equiv \tau \pmod n} a_x \, b_y \tag{21.1-1b}$$

The last equation may be rewritten as

$$h_\tau := \sum_{x=0}^{n-1} a_x \, b_{(\tau - x) \mod n} \tag{21.1-2}$$

That is, indices $\tau - x$ wrap around, it is a cyclic convolution. A convenient way to illustrate the cyclic convolution of two sequences is shown in figure 21.1-A.

21.1.1 Direct computation

A C++ implementation of the computation by definition is [FXT: `slow_convolution()` in convolution/slowcnvl.h]:

```
1  template <typename Type>
2  void slow_convolution(const Type *f, const Type *g, Type *h, ulong n)
3  // (cyclic) convolution:  h[] :=  f[] (*) g[]
```

```
          +--  0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
          |
          0:   0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
          1:   1  2  3  4   5  6  7  8   9 10 11 12  13 14 15  0
          2:   2  3  4  5   6  7  8  9  10 11 12 13  14 15  0  1
          3:   3  4  5  6   7  8  9 10  11 12 13 14  15  0  1  2

          4:   4  5  6  7   8  9 10 11  12 13 14 15   0  1  2  3
          5:   5  6  7  8   9 10 11 12  13 14 15  0   1  2  3  4
          6:   6  7  8  9  10 11 12 13  14 15  0  1   2  3  4  5
          7:   7  8  9 10  11 12 13 14  15  0  1  2   3  4  5  6

          8:   8  9 10 11  12 13 14 15   0  1  2  3   4  5  6  7
          9:   9 10 11 12  13 14 15  0   1  2  3  4   5  6  7  8
         10:  10 11 12 13  14 15  0  1   2  3  4  5   6  7  8  9
         11:  11 12 13 14  15  0  1  2   3  4  5  6   7  8  9 10

         12:  12 13 14 15   0  1  2  3   4  5  6  7   8  9 10 11
         13:  13 14 15  0   1  2  3  4   5  6  7  8   9 10 11 12
         14:  14 15  0  1   2  3  4  5   6  7  8  9  10 11 12 13
         15:  15  0  1  2   3  4  5  6   7  8  9 10  11 12 13 14


    +--   0   1   2   3   (a)
    |
    0:    0   1   2   4   ...
    1:    1   3   5 <--= h[5] contains a[2]*b[1]
    2:    4   8   9 <--= h[9] contains a[2]*b[2]
    3:    ...
  (b):
```

Figure 21.1-A: Semi-symbolic table of the cyclic convolution of two sequences (top). The entries denote where in the convolution the products of the input elements can be found (bottom).

```
4   // n := array length
5   {
6       for (ulong tau=0;  tau<n;  ++tau)
7       {
8           Type s = 0.0;
9           for (ulong k=0; k<n; ++k)
10          {
11              ulong k2 = tau - k;
12              if ( (long)k2<0 )  k2 += n;  // modulo n
13              s += (f[k]*g[k2]);
14          }
15          h[tau] = s;
16      }
17  }
```

The following version avoids the `if` statement in the inner loop:

```
1       for (ulong tau=0;  tau<n;  ++tau)
2       {
3           Type s = 0.0;
4           ulong k = 0;
5           for (ulong k2=tau;  k<=tau;  ++k, --k2)  s += (f[k]*g[k2]);
6           for (ulong k2=n-1;  k<n;     ++k, --k2)  s += (f[k]*g[k2]);  // wrapped around
7           h[tau] = s;
8       }
```

For length-n sequences this procedure involves proportional n^2 operations, therefore it is slow for large values of n. For short lengths the algorithm is just fine. Unrolled routines will offer good performance, especially for convolutions of fixed length. For medium length convolutions the splitting schemes given in section 27.2 on page 542 and section 38.2 on page 821 are applicable.

21.1.2 Computation via FFT

The Fourier transform provides us with an efficient way to compute convolutions that only uses proportional $n \log(n)$ operations. The *convolution property* of the Fourier transform is

$$\mathcal{F}\big[\, a \circledast b \,\big] \quad = \quad \mathcal{F}\big[\, a \,\big]\, \mathcal{F}\big[\, b \,\big] \tag{21.1-3}$$

That is, convolution in original space is element-wise multiplication in Fourier space. The statement can be motivated as follows:

$$\mathcal{F}\big[\,a\,\big]_k\,\mathcal{F}\big[\,b\,\big]_k \;=\; \sum_x a_x\,z^{k\,x}\sum_y b_y\,z^{k\,y} \tag{21.1-4a}$$

$$=\; \sum_x a_x\,z^{k\,x}\sum_{\tau-x} b_{\tau-x}\,z^{k\,(\tau-x)} \quad \text{where} \quad y=\tau-x \tag{21.1-4b}$$

$$=\; \sum_x\sum_{\tau-x} a_x\,z^{k\,x}b_{\tau-x}\,z^{k\,(\tau-x)} \;=\; \sum_\tau\left(\sum_x a_x\,b_{\tau-x}\right)z^{k\,\tau} \tag{21.1-4c}$$

$$=\; \left(\mathcal{F}\big[\sum_x a_x\,b_{\tau-x}\,\big]\right)_k \;=\; \big(\mathcal{F}\big[\,a\circledast b\,\big]\big)_k \tag{21.1-4d}$$

Rewriting relation 21.1-3 as

$$a\circledast b \;=\; \mathcal{F}^{-1}\big[\,\mathcal{F}\big[\,a\,\big]\,\mathcal{F}\big[\,b\,\big]\,\big] \tag{21.1-5}$$

tells us how to proceed. We give pseudo code for the cyclic convolution of two complex valued sequences `x[]` and `y[]`, result is returned in `y[]`:

```
1   procedure fft_cyclic_convolution(x[], y[], n)
2   {
3       complex x[0..n-1], y[0..n-1]
4
5       // transform data:
6       fft(x[], n, +1)
7       fft(y[], n, +1)
8
9       // convolution in transformed domain:
10      for i:=0 to n-1
11      {
12          y[i] := y[i] * x[i]
13      }
14
15       // transform back:
16      fft(y[], n, -1)
17
18      // normalize:
19      n1 := 1 / n
20      for i:=0 to n-1
21      {
22          y[i] := y[i] * n1
23      }
24  }
```

It is assumed that the procedure `fft()` does no normalization. For the normalization loop we precompute `1/n` and multiply as divisions are usually much slower than multiplications.

Relation 21.1-3 also holds for the more general z-transform (see section 21.5 on page 440). However, there is no (efficient) algorithm for the back-transform, so we cannot turn the relation

$$\mathcal{Z}\,[\,a\circledast b\,] \;=\; \mathcal{Z}\,[\,a\,]\;\mathcal{Z}\,[\,b\,] \tag{21.1-6}$$

into a practical algorithm for convolution.

21.1.3 Avoiding the revbin permutations

One can save the revbin permutations by observing that any DIF FFT is of the form

```
DIF_FFT_CORE(f, n);
revbin_permute(f, n);
```

and any DIT FFT is of the form

```
revbin_permute(f, n);
DIT_FFT_CORE(f, n);
```

Thereby a convolution routine that uses DIF FFTs for the forward transform and DIT FFTs as backward transform can omit the revbin permutations as demonstrated in the C++ implementation for the cyclic convolution of complex sequences [FXT: `fft_complex_convolution()` in convolution/fftcocnvl.cc]:

```
#define  DIT_FFT_CORE  fft_dit4_core_m1  // isign = -1
#define  DIF_FFT_CORE  fft_dif4_core_p1  // isign = +1
void
fft_complex_convolution(Complex * restrict f, Complex * restrict g,
                        ulong ldn, double v/*=0.0*/)
// (complex, cyclic) convolution:  g[] :=  f[] (*) g[]
// (use zero padded data for usual convolution)
// ldn := base-2 logarithm of the array length
// Supply a value for v for a normalization factor != 1/n
{
    const ulong n = (1UL<<ldn);

    DIF_FFT_CORE(f, ldn);
    DIF_FFT_CORE(g, ldn);
    if ( v==0.0 )  v = 1.0/n;
    for (ulong i=0; i<n; ++i)
    {
        Complex t = g[i] * f[i];
        g[i] = t * v;
    }
    DIT_FFT_CORE(g, ldn);
}
```

The signs of the two FFTs must be different but are else immaterial.

The so-called *auto convolution* (or *self convolution*) of a sequence is defined as the convolution of a sequence with itself: $h = a \circledast a$. The corresponding procedure needs only two instead of three FFTs.

21.1.4 Linear (acyclic) convolution

In the definition of the cyclic convolution (relations 21.1-1a and 21.1-1b) one can distinguish between those summands where the $x+y$ 'wrapped around' (i.e. $x+y = n+\tau$) and those where simply $x+y = \tau$ holds. These are (following the notation in [96]) denoted by $h^{(1)}$ and $h^{(0)}$ respectively. We have

$$h = h^{(0)} + h^{(1)} \qquad \text{where} \tag{21.1-7a}$$

$$h^{(0)} = \sum_{x \le \tau} a_x \, b_{\tau - x} \tag{21.1-7b}$$

$$h^{(1)} = \sum_{x > \tau} a_x \, b_{n+\tau - x} \tag{21.1-7c}$$

There is a simple way to separate $h^{(0)}$ and $h^{(1)}$ as the left and right half of a length-$2\,n$ sequence. This is just what the *linear convolution* (or *acyclic convolution*) does: Acyclic convolution of two (length-n) sequences a and b can be defined as that length-$2\,n$ sequence h which is the cyclic convolution of the *zero padded* sequences A and B:

$$A := [a_0, a_1, a_2, \ldots, a_{n-1}, 0, 0, \ldots, 0] \tag{21.1-8}$$

Same for B. Then the linear convolution is defined as

$$h = a \circledast_{lin} b \tag{21.1-9a}$$

$$h_\tau := \sum_{x=0}^{2\,n-1} A_x \, B_{\tau - x} \qquad \tau = 0, 1, 2, \ldots, 2\,n-1 \tag{21.1-9b}$$

As an illustration consider the convolution of the sequence $[1, 1, 1, 1]$ with itself: its linear self convolution is the length-8 sequence $[h_0][h_1] = [1, 2, 3, 4][3, 2, 1, 0]$, its cyclic self convolution is $[h_0 + h_1] = [4, 4, 4, 4]$.

```
 +--  0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
 |
 0:   0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
 1:   1  2  3  4   5  6  7  8   9 10 11 12  13 14 15 16
 2:   2  3  4  5   6  7  8  9  10 11 12 13  14 15 16 17
 3:   3  4  5  6   7  8  9 10  11 12 13 14  15 16 17 18

 4:   4  5  6  7   8  9 10 11  12 13 14 15  16 17 18 19
 5:   5  6  7  8   9 10 11 12  13 14 15 16  17 18 19 20
 6:   6  7  8  9  10 11 12 13  14 15 16 17  18 19 20 21
 7:   7  8  9 10  11 12 13 14  15 16 17 18  19 20 21 22

 8:   8  9 10 11  12 13 14 15  16 17 18 19  20 21 22 23
 9:   9 10 11 12  13 14 15 16  17 18 19 20  21 22 23 24
10:  10 11 12 13  14 15 16 17  18 19 20 21  22 23 24 25
11:  11 12 13 14  15 16 17 18  19 20 21 22  23 24 25 26

12:  12 13 14 15  16 17 18 19  20 21 22 23  24 25 26 27
13:  13 14 15 16  17 18 19 20  21 22 23 24  25 26 27 28
14:  14 15 16 17  18 19 20 21  22 23 24 25  26 27 28 29
15:  15 16 17 18  19 20 21 22  23 24 25 26  27 28 29 30
```

Figure 21.1-B: Semi-symbolic table for the (length-31) linear convolution of two length-16 sequences.

The semi-symbolic table for the acyclic convolution is given in figure 21.1-B. The elements in the lower right triangle do not 'wrap around' anymore, they go to extra buckets. Note there are 31 buckets labeled $0 \ldots 30$.

Linear convolution is *polynomial multiplication*: let $A = a_0 + a_1\,x + a_2\,x^2 + \ldots$, $B = b_0 + b_1\,x + b_2\,x^2 + \ldots$ and $C = A\,B = c_0 + c_1\,x + c_2\,x^2 + \ldots$ then

$$c_k \quad = \quad \sum_{i+j=k} a_i\,b_j \tag{21.1-10}$$

Chapter 27 on page 541 explains how fast convolution algorithms can be used for fast multiplication of multiprecision numbers.

The direct (slow) algorithm can be modified to compute just $h^{(0)}$ or $h^{(1)}$ [FXT: convolution/slowcnvlhalf.h]:

```
1   template <typename Type>
2   void slow_half_convolution(const Type *f, const Type *g, Type *h, ulong n, int h01)
3   // Half cyclic convolution.
4   // Part determined by h01 which must be 0 or 1.
5   // n := array length
6   {
7       if ( 0==h01 )  // compute h0:
8       {
9           for (ulong tau=0;  tau<n;  ++tau)
10          {
11              Type s0 = 0.0;
12              for (ulong k=0, k2=tau;  k<=tau;  ++k, --k2)  s0 += (f[k]*g[k2]);
13              h[tau] = s0;
14          }
15      }
16      else  // compute h1 (wrapped part):
17      {
18          for (ulong tau=0;  tau<n;  ++tau)
19          {
20              Type s1 = 0.0;
21              for (ulong k2=n-1, k=tau+1;  k<n;  ++k, --k2)  s1 += (f[k]*g[k2]);
22              h[tau] = s1;
23          }
24      }
25  }
```

The cost is half of what it is for the linear convolution. With the FFT-based computation of the convolution the parts $h^{(0)}$ and $h^{(1)}$ can be isolated using weighted convolution algorithms that are given in section 21.3 on page 435. The cost is that of a linear convolution. No algorithm to compute just one of $h^{(0)}$ or $h^{(1)}$ that is significantly cheaper is known.

21.2 Correlation

+--	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
|																
0:	0	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1:	1	0	15	14	13	12	11	10	9	8	7	6	5	4	3	2
2:	2	1	0	15	14	13	12	11	10	9	8	7	6	5	4	3
3:	3	2	1	0	15	14	13	12	11	10	9	8	7	6	5	4
4:	4	3	2	1	0	15	14	13	12	11	10	9	8	7	6	5
5:	5	4	3	2	1	0	15	14	13	12	11	10	9	8	7	6
6:	6	5	4	3	2	1	0	15	14	13	12	11	10	9	8	7
7:	7	6	5	4	3	2	1	0	15	14	13	12	11	10	9	8
8:	8	7	6	5	4	3	2	1	0	15	14	13	12	11	10	9
9:	9	8	7	6	5	4	3	2	1	0	15	14	13	12	11	10
10:	10	9	8	7	6	5	4	3	2	1	0	15	14	13	12	11
11:	11	10	9	8	7	6	5	4	3	2	1	0	15	14	13	12
12:	12	11	10	9	8	7	6	5	4	3	2	1	0	15	14	13
13:	13	12	11	10	9	8	7	6	5	4	3	2	1	0	15	14
14:	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	15
15:	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Figure 21.2-A: Semi-symbolic table for the (cyclic) correlation of two length-16 sequences.

+--	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
|																
0:	0	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17
1:	1	0	31	30	29	28	27	26	25	24	23	22	21	20	19	18
2:	2	1	0	31	30	29	28	27	26	25	24	23	22	21	20	19
3:	3	2	1	0	31	30	29	28	27	26	25	24	23	22	21	20
4:	4	3	2	1	0	31	30	29	28	27	26	25	24	23	22	21
5:	5	4	3	2	1	0	31	30	29	28	27	26	25	24	23	22
6:	6	5	4	3	2	1	0	31	30	29	28	27	26	25	24	23
7:	7	6	5	4	3	2	1	0	31	30	29	28	27	26	25	24
8:	8	7	6	5	4	3	2	1	0	31	30	29	28	27	26	25
9:	9	8	7	6	5	4	3	2	1	0	31	30	29	28	27	26
10:	10	9	8	7	6	5	4	3	2	1	0	31	30	29	28	27
11:	11	10	9	8	7	6	5	4	3	2	1	0	31	30	29	28
12:	12	11	10	9	8	7	6	5	4	3	2	1	0	31	30	29
13:	13	12	11	10	9	8	7	6	5	4	3	2	1	0	31	30
14:	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	31
15:	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Figure 21.2-B: Semi-symbolic table for the linear (acyclic) correlation of two length-16 sequences.

The *cyclic correlation* (or *circular correlation*) of two *real* length-n sequences $a = [a_0, a_1, \ldots, a_{n-1}]$ and $b = [b_0, b_1, \ldots, b_{n-1}]$ is defined as the length-n sequence h with elements h_τ as:

$$h_\tau \quad := \sum_{x-y \equiv \tau \bmod n} a_x \, b_y \tag{21.2-1}$$

The relation can be recast as

$$h_\tau \quad = \quad \sum_{x=0}^{n-1} a_x \, b_{(\tau+x) \bmod n} \tag{21.2-2}$$

The semi-symbolic table for the (cyclic) correlation is shown in figure 21.2-A. For the computation of the *linear* (or *acyclic*) correlation the sequences have to be zero-padded as in the algorithm for the linear convolution. The semi-symbolic table is shown in figure 21.2-B.

The *auto correlation* (or *self-correlation*) is the correlation of a sequence with itself, the correlation of two distinct sequences is also called *cross correlation*. The term *auto correlation function* (ACF) is often used for the auto correlation sequence.

21.2.1 Direct computation

A C++ implementation of the computation by the definition is [FXT: correlation/slowcorr.h]:

```
template <typename Type>
void slow_correlation(const Type *f, const Type *g, Type * restrict h, ulong n)
// Cyclic correlation of f[], g[], both real-valued sequences.
// n := array length
{
    for (ulong tau=0;  tau<n;  ++tau)
    {
        Type s = 0.0;
        for (ulong k=0; k<n; ++k)
        {
             ulong k2 = k + tau;
             if ( k2>=n )  k2 -= n;
             s += (g[k]*f[k2]);
        }
        h[tau] = s;
    }
}
```

The `if` statement in the inner loop is avoided by the following version:

```
    for (ulong tau=0;  tau<n;  ++tau)
    {
        Type s = 0.0;
        ulong k = 0;
        for (ulong k2=tau;  k2<n;  ++k, ++k2)  s += (g[k]*f[k2]);
        for (ulong k2=0;     k<n;  ++k, ++k2)  s += (g[k]*f[k2]);
        h[tau] = s;
    }
```

For the linear correlation one can avoid zero products:

```
template <typename Type>
void slow_correlation0(const Type *f, const Type *g, Type * restrict h, ulong n)
// Linear correlation of f[], g[], both real-valued sequences.
// n := array length
// Version for zero padded data:
//   f[k],g[k] == 0 for k=n/2 ... n-1
// n must be >=2
{
    const ulong nh = n/2;
    for (ulong tau=0; tau<nh; ++tau)  // k2 == tau + k
    {
        Type s = 0;
        for (ulong k=0, k2=tau;  k2<nh;  ++k, ++k2)  s += (f[k]*g[k2]);
        h[tau] = s;
    }

    for (ulong tau=nh; tau<n; ++tau)  // k2 == tau + k - n
    {
        Type s = 0;
        for (ulong k=n-tau, k2=0;  k<nh;  ++k, ++k2)  s += (f[k]*g[k2]);
        h[tau] = s;
    }
}
```

The algorithm involves proportional n^2 operations and is therefore slow with very long arrays.

21.2.2 Computation via FFT

A simple algorithm for fast correlation follows from the relation

$$h_\tau \quad = \quad \mathcal{F}^{-1}\big[\,\mathcal{F}\big[\,\overline{a}\,\big]\,\mathcal{F}\big[\,b\,\big]\,\big] \tag{21.2-3}$$

That is, use a convolution algorithm with one of the input sequences reversed (indices negated modulo n). For purely real sequences the relation is equivalent to complex conjugation of one of the inner transforms:

$$h_\tau \quad = \quad \mathcal{F}^{-1}\big[\,\mathcal{F}\big[\,a\,\big]^*\,\mathcal{F}\big[\,b\,\big]\,\big] \tag{21.2-4}$$

For the computation of self-correlation the latter relation is the only reasonable way to go: first transform the input sequence, then multiply each element by its complex conjugate and finally transform back. A C++ implementation is [FXT: correlation/fftcorr.cc]:

```
1   void
2   fft_correlation(double *f, double *g, ulong ldn)
3   // Cyclic correlation of f[], g[], both real-valued sequences.
4   // Result is written to g[].
5   // ldn := base-2 logarithm of the array length
6   {
7       const ulong n=(1UL<<ldn);
8       const ulong nh=(n>>1);
9
10      fht_real_complex_fft(f, ldn);  // real, imag part in lower, upper half
11      fht_real_complex_fft(g, ldn);
12
13      const double v = 1.0/n;
14      g[0]  *= f[0] * v;
15      g[nh] *= f[nh] * v;
16      for (ulong i=1,j=n-1; i<nh;  ++i,--j)  // real at index i, imag at index j
17      {
18          cmult_n(f[i], -f[j], g[i], g[j], v);
19      }
20
21      fht_complex_real_fft(g, ldn);
22  }
```

The function `cmult_n()` is defined in [FXT: aux0/cmult.h]:

```
1   static inline void
2   cmult_n(double c, double s, double &u, double &v, double dn)
3   // {u,v} <--| {dn*(u*c-v*s), dn*(u*s+v*c)}
4   { double t = u*s+v*c;  u *= c;  u -= v*s;  u *= dn;  v = t*dn; }
```

We note that relation 21.2-4 also holds for complex sequences.

21.2.3 Correlation and difference sets *

The linear auto-correlation of a sequence that contains zeros and ones only (a *delta sequence*) is the set of mutual differences of the positions of the ones, including multiplicity. An example:

```
[1, 1, 0, 1, 1, 0, 0, 0, 0, 0, 0]  <--= delta array  R
[4, 2, 1, 2, 1, 0, 0, 1, 2, 1, 2]  <--= linear auto correlation ACF
 0, 1, 2, 3, 4, 5,-5,-4,-3,-2,-1   <--= index
```

Element zero of the ACF tells us that there are four elements in R (each element has difference zero to just itself). Element one tells us that there are two pairs of consecutive elements, it is identical to the last element (element minus one). There is just one pair of elements in R whose indices differ by two (elements two and minus two of the ACF), and so on. Note that the ACF does not tell us where the elements with a certain difference are.

The delta array with ones at the 7 positions 0, 3, 4, 12, 18, 23, and 25 has the ACF

```
[7, 1, 1, 1, 1, 1, 1, 1, 1, 1, 0, 1, 1, 1, 1, 1, 0, 0, 1, 1, 1, 1, 1, 1, 0, 1, 0, (+symm.)]
 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10,       ...                                  26, <--= index
```

That is, a ruler of length 26 with marks only at the 7 given positions can be used to measure most of the distances up to 26 (the smallest missing distance is 10). Further, no distance appears more than once. Sequences with this property are called *Golomb rulers* and they are very hard to find.

If we allow for two rulers then the set of mutual differences in positions is the cross correlation. For this setting analogues of Golomb rulers (that do not have any missing differences) can be found. We use dots for zeros:

```
11..11..........11..11.....................................  <--= R1
1.1.....1.1.....1.....1.1.....1.1.....1.1.....1.1.....1.1  <--= R2
1111111111111111111111111111111111111111111111111111111111  <--= cross correlation
```

The rulers are binary representations of the evaluations $F(1/2)$ and $F(1/4)$ of a curious function given in section 36.10.1 on page 737.

21.3 Weighted Fourier transforms and convolutions

21.3.1 The weighted Fourier transform

We define a new kind of transform by slightly modifying the definition of the FT (formula 20.1-1a on page 393):

$$c = \mathcal{W}_v\,[a] \tag{21.3-1a}$$

$$c_k := \sum_{x=0}^{n-1} v_x\, a_x\, z^{x\,k} \qquad v_x \neq 0 \quad \forall x \tag{21.3-1b}$$

where $z := e^{\sigma\, 2\,\pi\, i/n}$. The sequence c shall be called (discrete) *weighted transform* of the sequence a with the weight (sequence) v. Note the v_x that entered: the weighted transform with $v_x = \frac{1}{\sqrt{n}}\ \forall x$ is just the usual Fourier transform. The inverse transform is

$$a = \mathcal{W}_v^{-1}\,[c] \tag{21.3-2a}$$

$$a_x = \frac{1}{n\, v_x} \sum_{k=0}^{n-1} c_k\, z^{-x\,k} \tag{21.3-2b}$$

This can be easily seen:

$$\mathcal{W}_v^{-1}\,[\mathcal{W}_v\,[a]]_y = \frac{1}{n\, v_y} \sum_{k=0}^{n-1} \sum_{x=0}^{n-1} v_x\, a_x\, z^{x\,k}\, z^{-y\,k} \tag{21.3-3a}$$

$$= \frac{1}{n} \sum_{k=0}^{n-1} \sum_{x=0}^{n-1} v_x\, \frac{1}{v_y}\, a_x\, z^{x\,k}\, z^{-y\,k} \tag{21.3-3b}$$

$$= \frac{1}{n} \sum_{x=0}^{n-1} v_x\, \frac{1}{v_y}\, a_x\, \delta_{x,y}\, n \quad = a_y \tag{21.3-3c}$$

Obviously all v_x have to be invertible. That $\mathcal{W}_v\left[\mathcal{W}_v^{-1}\,[a]\right]$ is also identity is apparent from the definitions.

Given an FFT routine it is trivial to set up a weighted Fourier transform. Pseudo code for the discrete weighted Fourier transform:

```
1  procedure weighted_ft(a[], v[], n, is)
2  {
3      for x:=0 to n-1
4      {
5          a[x] := a[x] * v[x]
6      }
7
8      fft(a[], n, is)
9  }
```

The inverse is essentially identical. Pseudo code for the inverse discrete weighted Fourier transform:

```
1  procedure inverse_weighted_ft(a[], v[], n, is)
2  {
3      fft(a[], n, -is)
4
5      for x:=0 to n-1
6      {
7          a[x] := a[x] / v[x]
8      }
9  }
```

The C++ implementations are given in [FXT: fft/weightedfft.cc].

21.3.2 Weighted convolution

Define the *weighted (cyclic) convolution* h_v by

$$h_v = a \circledast_{\{v\}} b \tag{21.3-4a}$$
$$= \mathcal{W}_v^{-1}\left[\mathcal{W}_v\left[a\right]\,\mathcal{W}_v\left[b\right]\right] \tag{21.3-4b}$$

Then, for the special case $v_x = V^x$, one has

$$h_v = h^{(0)} + V^n\, h^{(1)} \tag{21.3-5}$$

Here $h^{(0)}$ and $h^{(1)}$ are defined as in relation 21.1-7a on page 430. It is not hard to see why this is: up to the final division by the weight sequence, the weighted convolution is just the cyclic convolution of the two weighted sequences, which is for the element with index τ equal to

$$\sum_{x+y\equiv\tau \mod n} (a_x\, V^x)\,(b_y\, V^y) = \sum_{x\le\tau} a_x\, b_{\tau-x}\, V^\tau + \sum_{x>\tau} a_x\, b_{n+\tau-x}\, V^{n+\tau} \tag{21.3-6}$$

Final division of this element (by V^τ) gives $h^{(0)} + V^n\, h^{(1)}$ as stated.

The cases when V^n is some root of unity are particularly interesting: For $V^n = \pm i = \pm\sqrt{-1}$ one obtains the so-called *right-angle convolution*:

$$h_v = h^{(0)} \mp i\, h^{(1)} \tag{21.3-7}$$

This gives a nice possibility to directly use complex FFTs for the computation of a linear (acyclic) convolution of two real sequences: For length-n sequences the elements of the linear convolution with indices $0, 1, \ldots, n-1$ are the real part of the result, the elements $n, n+1, \ldots, 2\,n-1$ are the imaginary part. Choosing $V^n = -1$ leads to the *negacyclic convolution* (or *skew circular convolution*):

$$h_v = h^{(0)} - h^{(1)} \tag{21.3-8}$$

Cyclic, negacyclic and right-angle convolution can be understood as polynomial products modulo the polynomials $z^n - 1$, $z^n + 1$ and $z^n \pm i$, respectively (see [207]).

C++ implementations of the weighted-, negacyclic- and right-angle (self) convolution are given in [FXT: convolution/weightedconv.cc].

+--	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
|																
0:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	0-
2:	2	3	4	5	6	7	8	9	10	11	12	13	14	15	0-	1-
3:	3	4	5	6	7	8	9	10	11	12	13	14	15	0-	1-	2-
4:	4	5	6	7	8	9	10	11	12	13	14	15	0-	1-	2-	3-
5:	5	6	7	8	9	10	11	12	13	14	15	0-	1-	2-	3-	4-
6:	6	7	8	9	10	11	12	13	14	15	0-	1-	2-	3-	4-	5-
7:	7	8	9	10	11	12	13	14	15	0-	1-	2-	3-	4-	5-	6-
8:	8	9	10	11	12	13	14	15	0-	1-	2-	3-	4-	5-	6-	7-
9:	9	10	11	12	13	14	15	0-	1-	2-	3-	4-	5-	6-	7-	8-
10:	10	11	12	13	14	15	0-	1-	2-	3-	4-	5-	6-	7-	8-	9-
11:	11	12	13	14	15	0-	1-	2-	3-	4-	5-	6-	7-	8-	9-	10-
12:	12	13	14	15	0-	1-	2-	3-	4-	5-	6-	7-	8-	9-	10-	11-
13:	13	14	15	0-	1-	2-	3-	4-	5-	6-	7-	8-	9-	10-	11-	12-
14:	14	15	0-	1-	2-	3-	4-	5-	6-	7-	8-	9-	10-	11-	12-	13-
15:	15	0-	1-	2-	3-	4-	5-	6-	7-	8-	9-	10-	11-	12-	13-	14-

Figure 21.3-A: Semi-symbolic table for the negacyclic convolution. The products that enter with negative sign are indicated with a postfix minus at the corresponding entry.

The semi-symbolic table for the negacyclic convolution is shown in figure 21.3-A. With right-angle convolution the minuses have to be replaced by $i = \sqrt{-1}$ which means the wrap-around (i.e. $h^{(1)}$) elements

go to the imaginary part. With real input one thereby effectively separates $h^{(0)}$ and $h^{(1)}$. Thereby the linear convolution of real sequences can be computed using the complex right-angle convolution.

With routines for both cyclic and negacyclic convolution the parts $h^{(0)}$ and $h^{(1)}$ can be computed as sum and difference, respectively. Thereby all expressions of the form $\alpha\, h^{(0)} + \beta\, h^{(1)}$ where $\alpha, \beta \in \mathbb{C}$ can be computed.

The direct (slow, proportional n^2) computation can be obtained by a minimal modification of the non-weighted convolution algorithm [FXT: convolution/slowweightedcnvl.h]:

```
1   template <typename Type>
2   void slow_weighted_convolution(const Type *f, const Type *g, Type *h, ulong n, Type w)
3   // weighted (cyclic) convolution:  h[] :=  f[] (*)_w g[]
4   // n := array length
5   {
6       for (ulong tau=0;  tau<n;  ++tau)
7       {
8           ulong k = 0;
9           Type s0 = 0.0;
10          for (ulong k2=tau;  k<=tau;  ++k, --k2)  s0 += (f[k]*g[k2]);
11          Type s1 = 0.0;
12          for (ulong k2=n-1;  k<n;     ++k, --k2)  s1 += (f[k]*g[k2]);  // wrapped around
13          h[tau] = s0 + s1*w;
14      }
15  }
```

21.4 Convolution using the MFA

We give an algorithm for convolution that use the matrix Fourier algorithm (MFA, see section 20.11 on page 424). The MFA is used for the forward transform and the transposed algorithm (TMFA) for the backward transform. The elements of each row are assumed to lie contiguous in memory. For the sake of simplicity auto convolution is considered. The matrix FFT convolution algorithm:

1. Apply a (length R) FFT on each column.
   (*memory access with C-skips*)
2. Multiply each matrix element (index r, c) by $\exp(+\sigma\, 2\,\pi\, i\, r\, c/n)$.
3. Apply a (length C) FFT on each row.
   (*memory access without skips*)
4. Complex square row (element-wise).
5. Apply a (length C) FFT on each row (of the transposed matrix).
   (*memory access is without skips*)
6. Multiply each matrix element (index r, c) by $\exp(-\sigma\, 2\,\pi\, i\, r\, c/n)$.
7. Apply a (length R) FFT on each column (of the transposed matrix).
   (*memory access with C-skips*)

Note that steps 3, 4, and 5 constitute a length-C convolution on each row.

C++ implementations of the cyclic and linear convolution using the described algorithm are given in [FXT: convolution/matrixfftcnvl.cc]. The code for self-convolution can be found in [FXT: convolution/matrixfftcnvla.cc].

With the weighted convolutions in mind we reformulate the matrix (self-) convolution algorithm given on page 437:

1. Apply a FFT on each column.

2. On each row apply the weighted convolution with $V^C = e^{2\pi i r/R} = 1^{r/R}$ where R is the total number of rows, $r = 0..R-1$ the index of the row, C the length of each row (equivalently, the total number columns)
3. Apply a FFT on each column (of the transposed matrix).

We first consider the special cases of two and three rows and then formulate an MFA-based algorithm for the convolution of real sequences.

The case $R = 2$

Define s and d as the sums and differences of the lower and higher halves of a given sequence x:

$$s := x^{(0/2)} + x^{(1/2)} \tag{21.4-1a}$$
$$d := x^{(0/2)} - x^{(1/2)} \tag{21.4-1b}$$

Then the cyclic auto convolution of the sequence x can be obtained by two half-length convolutions of s and d as

$$x \circledast x = \frac{1}{2}\left[s \circledast s + d \circledast_- d, \quad s \circledast s - d \circledast_- d\right] \tag{21.4-2}$$

where the symbols $\circledast$ and $\circledast_-$ stand for cyclic and negacyclic convolution, respectively (see section 21.3 on page 435). The equivalent formula for the cyclic convolution of two sequences x and y is

$$x \circledast y = \frac{1}{2}\left[s_x \circledast s_y + d_x \circledast_- d_y, \quad s_x \circledast s_y - d_x \circledast_- d_y\right] \tag{21.4-3}$$

where

$$s_x := x^{(0/2)} + x^{(1/2)} \tag{21.4-4a}$$
$$d_x := x^{(0/2)} - x^{(1/2)} \tag{21.4-4b}$$
$$s_y := y^{(0/2)} + y^{(1/2)} \tag{21.4-4c}$$
$$d_y := y^{(0/2)} - y^{(1/2)} \tag{21.4-4d}$$

Now use the fact that an linear convolution is computed by a cyclic convolution of zero-padded sequences whose upper halves are simply zero, so $s_x = d_x = x$ and $s_y = d_y = y$. Then relation 21.4-3 reads:

$$x \circledast_{lin} y = \frac{1}{2}\left[x \circledast y + x \circledast_- y, \quad x \circledast y - x \circledast_- y\right] \tag{21.4-5}$$

And for the acyclic auto convolution:

$$x \circledast_{lin} x = \frac{1}{2}\left[x \circledast x + x \circledast_- x, \quad x \circledast x - x \circledast_- x\right] \tag{21.4-6}$$

The lower and upper halves of the linear convolution can be obtained from the sum and difference of the cyclic and the negacyclic convolution.

The case $R = 3$

Let $\omega = \frac{1}{2}(1 + i\sqrt{3})$ and define

$$A := x^{(0/3)} + x^{(1/3)} + x^{(2/3)} \tag{21.4-7a}$$
$$B := x^{(0/3)} + \omega\, x^{(1/3)} + \omega^2 x^{(2/3)} \tag{21.4-7b}$$
$$C := x^{(0/3)} + \omega^2 x^{(1/3)} + \omega\, x^{(2/3)} \tag{21.4-7c}$$

Then, if $h := x \circledast x$, one has

$$h^{(0/3)} = A \circledast A + \quad B \circledast_{\{\omega\}} B + \quad C \circledast_{\{\omega^2\}} C \tag{21.4-8a}$$

$$h^{(1/3)} = A \circledast A + \omega^2 B \circledast_{\{\omega\}} B + \omega \; C \circledast_{\{\omega^2\}} C \tag{21.4-8b}$$

$$h^{(2/3)} = A \circledast A + \omega \; B \circledast_{\{\omega\}} B + \omega^2 C \circledast_{\{\omega^2\}} C \tag{21.4-8c}$$

For real valued data C is the complex conjugate (*cc.*) of B and (with $\omega^2 = cc.\omega$) $B \circledast_{\{\omega\}} B$ is the *cc.* of $C \circledast_{\{\omega^2\}} C$ and therefore every $B \circledast_{\{.\}} B$-term is the cc. of the $C \circledast_{\{.\}} C$-term in the same line. Is there a nice and general scheme for real valued convolutions based on the MFA? Read on for the positive answer.

21.4.1 Convolution of real valued data using the MFA

Consider the MFA-algorithm for the cyclic convolution as given on page 437 but with real input data: for row 0 which is real after the column FFTs one needs to compute the usual cyclic convolution; for row $R/2$ which is also purely real after the column FFTs a negacyclic convolution is needed[1], the code for negacyclic convolution is given on page 515.

All other weighted convolutions involve complex computations, but it is easy to see how to reduce the work by 50 percent: as the result must be real the data in row number $R - r$ must, because of the symmetries of the real and imaginary part of the (inverse) Fourier transform of real data, be the complex conjugate of the data in row r. Therefore one can use real FFTs (R2CFTs) for all column-transforms for step 1 and half-complex to real FFTs (C2RFTs) for step 3.

Let the computational cost of a cyclic (real) convolution be q, then

- For even values of R one must perform one cyclic (row 0), one negacyclic (row $R/2$) and $R/2 - 2$ complex weighted convolutions (rows $1, 2, \ldots, R/2 - 1$)
- For R odd one must perform 1 cyclic (row 0) and $(R - 1)/2$ complex weighted convolutions (rows $1, 2, \ldots, (R - 1)/2$)

Now assume, slightly simplifying, that the cyclic and the negacyclic real convolution involve the same number of computations and that the cost of a weighted complex convolution is twice that. Then in both cases above the total work is exactly half of that for the complex case, which is what one expects from a real world real valued convolution algorithm.

For the computation of the linear convolution one can use the right angle convolution (and complex FFTs in the column passes), see section 21.3 on page 435.

21.4.2 Mass storage convolution using the MFA

Algorithms on data sets that do not fit into physical RAM are sometimes called *external* or *out of core* algorithms. Simply using the virtual memory mechanism of the operating system is not an option, eternal hard disk activity would be the consequence in most cases. We give a method for the *mass storage convolution*. It is based on the matrix FFT convolution algorithm given on page 437. The number of disk seeks has to be kept minimal because these are slow operations which degrade performance unacceptably if they occur too often,

The crucial modification of the use of the MFA is to *not* choose R and C as close as possible to $\sqrt{n}$ as is usually done. Instead one chooses R to be minimal so that the row length C corresponds to the biggest data set that fits into the available RAM. We now analyze how the number of seeks depends on the choice of R and C: In what follows it is assumed that the data lies in memory as row_0, row_1, ..., row_{R-1}. In other words, the elements of each rows lie contiguous in memory. Further let $\alpha \geq 2$ be the number of times the data set exceeds the available RAM size.

[1] For odd values of R there is no such row and no negacyclic convolution is needed.

In step 1 and 3 of the convolution algorithm given on page 437 one reads from disk (row by row, involving R seeks) the number of columns that just fit into RAM, does the (many, short) column-FFTs, writes back (again R seeks), and proceeds to the next block; this happens for α of these blocks, giving a total of $4\,\alpha\,R$ seeks for steps 1 and 3.

In step 2 one reads (α times) blocks of one or more rows, which lie in contiguous portions of the disk, perform the FFT on the rows, and write back to disk. This gives a total of $2\,\alpha$ seeks for step 2.

Thereby there are $2\,\alpha + 4\,\alpha\,R$ seeks during the whole computation, which is minimized by the choice of maximal C. This means that one chooses a shape of the matrix so that the rows are as big as possible subject to the constraint that they have to fit into main memory, which in turn means there are $R = \alpha$ rows, leading to an optimal seek count of $K = 2\,\alpha + 4\,\alpha^2$.

Let S_D be the seek time of the hard disk in seconds, W_D be the disk transfer rate in megabyte per second. Further let A be the amount of available RAM in megabytes. Then the total time spent for disk operations is the sums of the time spent in seeks and the time for reading and writing:

$$T = K\,S_D + 6\,\frac{\alpha\,A}{W_D} \tag{21.4-9a}$$

$$= \left(2\,\alpha + 4\,\alpha^2\right)\,S_D + 6\,\frac{\alpha\,A}{R_D} \tag{21.4-9b}$$

We give two examples for an FFT whose size that exceeds the available RAM by a factor of $\alpha = 16$ (so the number of seeks is $k = 1056$).

With a machine where the disk seek takes 10 milliseconds ($S_D = 0.010$) about 10 seconds are needed for the seeks. Further assume we have $A = 64$ Megabytes of RAM available so the transform size is 1 GB corresponding to 128 million (double precision) floats. The disk transfer rate shall be $W_D = 10$ MB/sec. Then the overhead for the read and write would amount to $6 \cdot 1024MB/(10MB/sec) \approx 615sec$ or approximately 10 minutes. So the total disk activity takes approximately 10.5 minutes.

With a workstation machine with $S_D = 0.005$ sec (5 milliseconds), $W_d = 50$ MB/sec, and 512 MB of RAM: the transform size is 8 GB (1 G doubles) and we obtain $T = 5.28{+}983.04 = 988.32$ or approximately 16.5 minutes.

In both cases the CPU will be busy for several minutes which is in the same order as the time for the disk activity. Therefore, when multi-threading is available, one may want to use a *double buffer* variant: Choose the row length so that it fits twice into the RAM workspace; then let always one (CPU-intensive) thread do the FFTs in one of the scratch spaces and the other (hard disk intensive) thread write back the data from the other scratch space and read the next data to be processed. With not too small main memory (and not too slow hard disk) and some fine tuning double buffering can to keep the CPU busy during much of the hard disk operations. Thereby most of the disk time can be 'hidden' as the CPU activity does not stall waiting for data.

The mass storage convolution as described was used for the calculation of the number

$$9^{9^9} \approx 0.4281247 \cdot 10^{369,693,100} \tag{21.4-10}$$

on a 32-bit machine in 1999. The computation used two files of size 2 Gigabytes each and took less than eight hours on a system with a AMD K6/2 CPU at 366 MHz with 66 MHz memory. The log-file of the computation is [hfloat: examples/run1-pow999.txt]. The computation did not use the double-buffer technique.

21.5 The z-transform (ZT)

In this section we will learn a technique to compute the Fourier transform by a cyclic convolution. In fact, the transform computed is the z-transform, a more general transform that in a special case is identical

to the Fourier transform.

The discrete z-transform (ZT) of a length-n sequence a is a length-n sequence c defined by

$$c = \mathcal{Z}[a] \tag{21.5-1a}$$

$$c_k := \sum_{x=0}^{n-1} a_x \, z^{k\,x} \tag{21.5-1b}$$

The z-transform is a linear transformation. It is not an orthogonal transformation unless z is a root of unity. For $z = e^{\pm 2\,\pi\,i/n}$ the z-transform specializes to the discrete Fourier transform. An important property is the convolution property:

$$\mathcal{Z}[a \circledast b] = \mathcal{Z}[a]\ \mathcal{Z}[b] \tag{21.5-2}$$

Convolution in original space corresponds to ordinary (element-wise) multiplication in z-space. This can be turned into an efficient convolution algorithm for the special case of the Fourier transform but not in general because no efficient algorithm for the inverse transform is known.

21.5.1 Computation via convolution (Bluestein's algorithm)

Noting that

$$x\,k = \frac{1}{2}\left(x^2 + k^2 - (k-x)^2\right) \tag{21.5-3}$$

we find, for element c_k of the Fourier transform of the sequence a,

$$c_k = \sum_{x=0}^{n-1} a_x \, z^{x\,k} = z^{k^2/2} \left[\sum_{x=0}^{n-1} \left(a_x \, z^{x^2/2}\right)\, z^{-(k-x)^2/2}\right] \tag{21.5-4}$$

The expression in brackets is a cyclic convolution of the sequence $a_x \, z^{x^2/2}$ with the sequence $z^{-x^2/2}$.

This leads to the algorithm for the *chirp z-transform*:

1. Multiply the sequence a element-wise with $z^{x^2/2}$.
2. Convolve the resulting sequence with the sequence $z^{-x^2/2}$.
3. Multiply element-wise with the sequence $z^{k^2/2}$.

The above algorithm constitutes a fast algorithm for the ZT because fast convolution is possible via FFT. The idea is due to Bluestein [46], a detailed description of the algorithm for computing the Bluestein FFT is given in [250].

21.5.2 Arbitrary length FFT by ZT

The length n of the input sequence a for the fast z-transform is not limited to highly composite values: For values of n where a FFT is not feasible pad the sequence with zeros up to a length L with $L >= 2\,n$ such that a length-L FFT can be computed (highly composite L, for example a power of two).

As the Fourier transform is the special case $z = e^{\pm 2\,\pi\,i/n}$ of the ZT the chirp-ZT algorithm constitutes an FFT algorithm for sequences of arbitrary length.

The transform takes a few times more than an direct FFT. The worst case (if only FFTs for n a power of 2 are available) is $n = 2^p + 1$: one must perform three FFTs of length $L = 2^{p+1} \approx 2\,n$ for the computation of the convolution. So the total work amounts to about 6 times the work a FFT of length n. It is possible to lower this 'worst case factor' to 3 by using highly composite L slightly greater than n.

For multiple computations of z-transforms of the same length one may want to store the Fourier transform of the sequence $z^{k^2/2}$ as it does not change. Thereby the worst case is reduced to a factor 2 with highly composite FFTs and 4 if FFTs are available for powers of two only.

A C++ implementation of the Fourier transform for sequences of arbitrary length is given in [FXT: chirpzt/fftarblen.cc]:

```
void
fft_arblen(Complex *x, ulong n, int is)
// Arbitrary length FFT.
{
    const ulong ldnn = 1 + ld( (n << 1) - 1 );
    const ulong nn = (1UL<<ldnn);  // smallest power of 2 >= 2*n

    Complex *f = new Complex[nn];
    copy(x, f, n);
    null(f+n, nn-n);

    Complex *w = new Complex[nn];
    make_fft_chirp(w, n, nn, is);
    multiply(f, n, w);

    double *dw = (double *)w;
    for (ulong k=1; k<2*n; k+=2)  dw[k] = -dw[k];  // =^= make_fft_chirp(w, n, nn, -is);

    fft_complex_convolution(w, f, ldnn);

    if ( n & 1 )  subtract(f, n, f+n);  // odd n:  negacyclic convolution
    else          add(f, n, f+n);       // even n:  cyclic convolution

    make_fft_chirp(w, n, nn, is);
    multiply(w, n, f);

    copy(w, x, n);
    delete [] w;
    delete [] f;
}
```

The auxiliary routine `make_fft_chirp()` is

```
static inline void
make_fft_chirp(Complex *w, ulong n, ulong nn, int is)
// For k=0..n-1:   w[k] := exp( is * k*k * (i*2*PI/n)/2 )
//   where i = sqrt(-1)
// For k=n..nn-1:  w[k] = 0
{
    double phi = 1.0*is*M_PI/n;  // == (i*2*Pi/n)/2
    ulong k2 = 0,  n2 = 2*n;
    for (ulong k=0; k<n; ++k)
    {
        w[k] = SinCos(phi*k2);
        k2 += (2*k+1);
        if ( k2>n2 )  k2 -= n2;
        // here:  k2 == (k*k) mod 2*n;
    }
    null(w+n, nn-n);
}
```

21.5.3 Fractional Fourier transform by ZT

The z-transform with $z = e^{\alpha\, 2\,\pi\, i/n}$ is called the *fractional Fourier transform* in [24]. We note that the term is usually used for the fractional order transform given as relation 24.14-6 on page 522, see also [215, ch.13].

For $\alpha = \pm 1$ one again obtains the usual Fourier transform. The fractional Fourier transform can be used for the computation of the Fourier transform of sequences with only few nonzero elements and for the exact detection of frequencies that are not integer multiples of the lowest frequency of the DFT.

A C++ implementation of the fractional Fourier transform for sequences of arbitrary length is given in [FXT: chirpzt/fftfract.cc]:

```
1   void
2   fft_fract(Complex *x, ulong n, double v)
3   // Fractional (fast) Fourier transform.
4   {
5       const ulong ldnn = 1 + ld( (n << 1) - 1 );
6       const ulong nn = (1UL<<ldnn);  // smallest power of 2 >= 2*n
7
8       Complex *f = new Complex[nn];
9       copy(x, f, n);
10      null(f+n, nn-n);
11
12      Complex *w = new Complex[nn];
13      make_fft_fract_chirp(w, v, n, nn);
14
15      for (ulong j=0; j<n; ++j)  f[j] *= w[j];
16      for (ulong j=0; j<nn; ++j)  w[j] = conj(w[j]);
17
18      fft_complex_convolution(w, f, ldnn);
19
20      make_fft_fract_chirp(w, v, n, nn);
21
22      for (ulong j=0; j<n; ++j)  w[j] *= f[j];
23
24      copy(w+n, x, n);
25      delete [] w;
26      delete [] f;
27  }
```

The auxiliary routine `make_fft_fract_chirp()` is

```
1   static inline void
2   make_fft_fract_chirp(Complex *w, double v, ulong n, ulong nn)
3   // For k=0..nn-1:   w[k] == exp(v*sqrt(-1)*k*k*2*pi*/n/2)
4   {
5       const double phi = v*2.0*M_PI/n/2;
6       ulong n2 = 2*n;
7       ulong np=0;
8       for (ulong k=0; k<nn; ++k)
9       {
10          w[k] = SinCos(phi*np);
11
12          np += ((k<<1)+1);  // np == (k*k)%n2
13          if ( np>=n2 )  np -= n2;
14      }
15  }
```

21.6 Prime length FFTs

For the computation of FFTs for sequences whose length is prime we can exploit the existence of primitive roots. We will be able to express the transform of all but the first element as a cyclic convolution of two sequences whose length is reduced by one.

Let p be prime, then an element g exists so that the least positive exponent e so that $g^e \equiv 1 \bmod p$ is $e = p - 1$. The element g is called a *generator* (or *primitive root*) modulo p (see section 37.5 on page 762). Every nonzero element modulo p can uniquely be expressed as a power g^e where $0 \le e < p-1$. For example, a generator modulo $p = 11$ is $g = 2$, its powers are:

$$g^0 \equiv 1 \quad g^1 \equiv 2 \quad g^2 \equiv 4 \quad g^3 \equiv 8 \quad g^4 \equiv 5 \quad g^5 \equiv 10 \equiv -1 \quad g^6 \equiv 9 \quad g^7 \equiv 7 \quad g^8 \equiv 3 \quad g^9 \equiv 6 \quad g^{p-1} \equiv 1$$

Likewise, we can express any nonzero element as a negative power of g. Let $h = g^{-1}$, then with our example, $h \equiv 6$ and

$$h^0 \equiv 1 \quad h^1 \equiv 6 \quad h^2 \equiv 3 \quad h^3 \equiv 7 \quad h^4 \equiv 9 \quad h^5 \equiv 10 \equiv -1 \quad h^6 \equiv 5 \quad h^7 \equiv 8 \quad h^8 \equiv 4 \quad h^9 \equiv 2 \quad h^{p-1} \equiv 1$$

This is just the reversed sequence of values. Let the C be the Fourier transform of length-p sequence A:

$$C_k \quad = \quad \sum_{x=0}^{p-1} A_x \, W^{\sigma \, x \, k} \tag{21.6-1}$$

where $W = \exp(2\,i\,\pi/p)$ and $\sigma = \pm 1$ is the sign of the transform. We split the computation of the Fourier transform in two parts, we compute the first element of the transform as

$$C_0 \;=\; \sum_{x=0}^{p-1} A_x \tag{21.6-2}$$

Now it remains to compute, for $1 \le k \le p-1$,

$$C_k \;=\; A_0 + \sum_{x=1}^{p-1} A_x\, W^{\sigma\, x\, k} \tag{21.6-3}$$

Note the lower index of the sum. We write $k \equiv g^e$ and $x \equiv g^{-f}$ (modulo p), so

$$C_{(g^e)} - A_0 \;=\; \sum_{f=0}^{p-2} A_{(g^{-f})}\, W^{\sigma\,(g^{-f})\,(g^e)} \;=\; \sum_{f=0}^{p-2} A_{(g^{-f})}\, W^{\sigma\,(g^{e-f})} \tag{21.6-4}$$

The sum is a cyclic convolution of the sequences $W^{(g^w)}$ and $A_{(g^{-w})}$ where $0 \le w \le p-2$. The main algorithm (ignoring the constant terms A_0 and C_0) can be outlined as follows:

1. Compute A^* and W^* by permuting the sequences A and W.
2. Compute C^* the cyclic convolution of A^* and W^*.
3. Compute W by permuting W^*.

The method is given in [218], it is called *Rader's algorithm*. We implement it in pari/gp:

```
ft_rader(a, is=+1)=
\\ Fourier transform for prime lengths (Rader's algorithm)
{
    local(n, a0, c0, g, w);
    local(c, ixp, ixm, pa, pw, t);
    n = length(a);

    a0 = a[1];  c0 = sum(j=1, n, a[j]);  \\ constant terms

    \\ prepare permutations:
    g = znprimroot(n);  ixp = vector(n, j, component( g^(j-1), 2) );
    g = g^(-1);  ixm = vector(n, j, component( g^(j-1), 2) );

    \\ permute sequence W:
    w = is*2*I*Pi/n;  pw = vector(n-1, j, exp(w*ixp[j]) );

    \\ permute sequence A:
    pa = vector(n-1);  for (j=1, n-1, pa[j]=a[1+ixp[1+n-j]] );

    \\ cyclic convolution of permuted sequences:
    t = cconv(pa, pw);  \\ cyclic convolution

    \\ set C_0, and add A_0 to each C_k:
    c = vector(n);  c[1] = c0;  for (k=1, n-1, c[1+k]=t[k]+a0);

    \\ permute to obtain result:
    t = vector(n);  t[1] = c[1];  for (k=2, n, t[1+ixp[k-1]]=c[k]);
    return( t );
}
```

With a (slow) implementation of the cyclic convolution and DFT we can check whether the method works by comparing the results:

```
cconv(a, b)=
/* Cyclic convolution (direct computation, n^2 operations) */
/* Example: cconv([a,b],[c,d]) ==> [b*d + c*a, a*d + c*b] */
{
    local(n, f, s, k, k2);
    n = length(a);
    f = vector(n);
```

```
8       for (tau=0, n-1,  \\ tau = k + k2
9           s0 = 0;  k = 0;  k2 = tau;
10          while (k<=tau, s0 += (a[k+1]*b[k2+1]);  k++; k2--);
11          s1 = 0;  k2 = n-1; \\ k=tau+1
12          while (k<n,    s1 += (a[k+1]*b[k2+1]);  k++; k2--);
13          f[tau+1] = s0 + s1;
14      );
15      return( f );
16  }
```

```
1   dft(a, is=+1)=
2   /* Complex Fourier transform (direct computation, n^2 operations) */
3   {
4       local(n, f, s, ph0, ph);
5       n = length(a);
6       f = vector(n);
7       ph0 = is*2*Pi*I/n;
8       for (k=0, n-1,
9           ph = ph0 * k;
10          f[k+1] = sum (x=0, n-1,  a[x+1] * exp(ph*x) );
11      );
12      return( f );
13  }
```

In order to turn the algorithm into a fast Fourier transform we need to compute the convolution via fast transforms of length $(p-1)$. This is trivially possible when $p-1=2^q$, for example when $p=5$ or $p=17$. As $p-1$ is always divisible by two, we can split at least once. For $p=11$ we have $(p-1)/2=5$ so we can again use Rader's algorithm and length-4 transforms.

The method can be used to generate code for short (prime) length FFTs. One should precompute the permuted and transformed sequence of the powers of the primitive root W. Thereby only two FFTs of length $(p-1)$ will be needed for a length-p transform.

Chapter 22

The Walsh transform and its relatives

We describe several variants of the *Walsh transform*, sometimes called *Walsh-Hadamard transform* or just *Hadamard transform*. The Walsh transform has the same complexity as the Fourier transform but does not involve any multiplications. In fact, one can obtain a Walsh transform routine by removing all multiplications (with sines and cosines) in a given FFT routine.

We also give related transforms like the slant transform and the Reed-Muller transform. The dyadic convolution that can be computed efficiently by the Walsh transform is introduced.

22.1 The Walsh transform: Walsh-Kronecker basis

How to make a Walsh transform out of your FFT:
'*Replace exp*(*anything*) *by* 1*, done.*'

Removing all *exp*(anything) from the radix-2, decimation in time Fourier transform we obtain

```
 1  void slow_walsh_wak_dit2(double *f, ulong ldn)
 2  // (this routine has a problem)
 3  {
 4      ulong n = (1UL<<ldn);
 5      for (ulong ldm=1; ldm<=ldn; ++ldm)
 6      {
 7          const ulong m = (1<<ldm);
 8          const ulong mh = (m>>1);
 9          for (ulong j=0; j<mh; ++j)
10          {
11              for (ulong r=0; r<n; r+=m)
12              {
13                  const ulong t1 = r+j;
14                  const ulong t2 = t1+mh;
15                  double u = f[t1];
16                  double v = f[t2];
17                  f[t1] = u+v;
18                  f[t2] = u-v;
19              }
20          }
21      }
22  }
```

The transform involves proportional $n \log_2(n)$ additions (and subtractions) and no multiplication at all. The transform, as given, is its own inverse up to a factor $1/n$. The Walsh transform of integer input is integral.

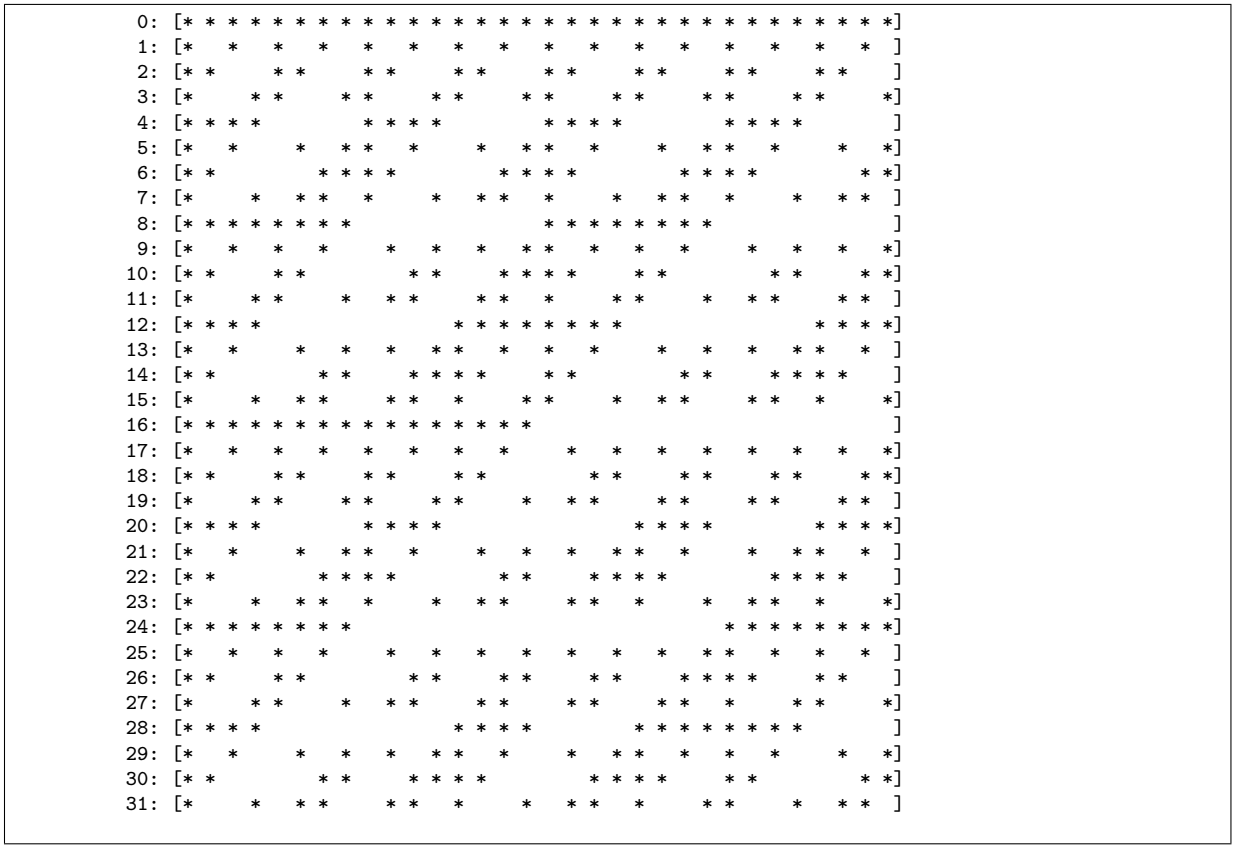

Figure 22.1-A: Basis functions for the Walsh transform (Walsh-Kronecker basis). Asterisks denote the value $+1$, blank entries denote -1.

As the `slow` in the name shall suggest, the implementation has a problem as given. The memory access pattern is highly non-local. Let's make a slight improvement: here we just took the radix-2 DIT FFT code from section 20.3.1.3 on page 399 and threw away all trigonometric computations (and multiplications). But the swapping of the inner loops, that we did for the FFT in order to save trigonometric computations is now of no advantage anymore. So we try the following [FXT: `walsh_wak_dit2()` in walsh/walshwak2.h]:

```
1   template <typename Type>
2   void walsh_wak_dit2(Type *f, ulong ldn)
3   // Transform wrt. to Walsh-Kronecker basis (wak-functions).
4   // Radix-2 decimation in frequency (DIF) algorithm.
5   {
6       ulong n = (1UL<<ldn);
7       for (ulong ldm=1; ldm<=ldn; ++ldm)
8       {
9           const ulong m = (1UL<<ldm);
10          const ulong mh = (m>>1);
11          for (ulong r=0; r<n; r+=m)
12          {
13              ulong t1 = r;
14              ulong t2 = r+mh;
15              for (ulong j=0;  j<mh;  ++j, ++t1, ++t2)
16              {
17                  Type u = f[t1];
18                  Type v = f[t2];
19                  f[t1] = u + v;
20                  f[t2] = u - v;
21              }
22          }
23      }
24  }
```

The performance impact is quite drastic. For $n = 2^{21}$ (and type double, 16 MByte of memory) it gives a

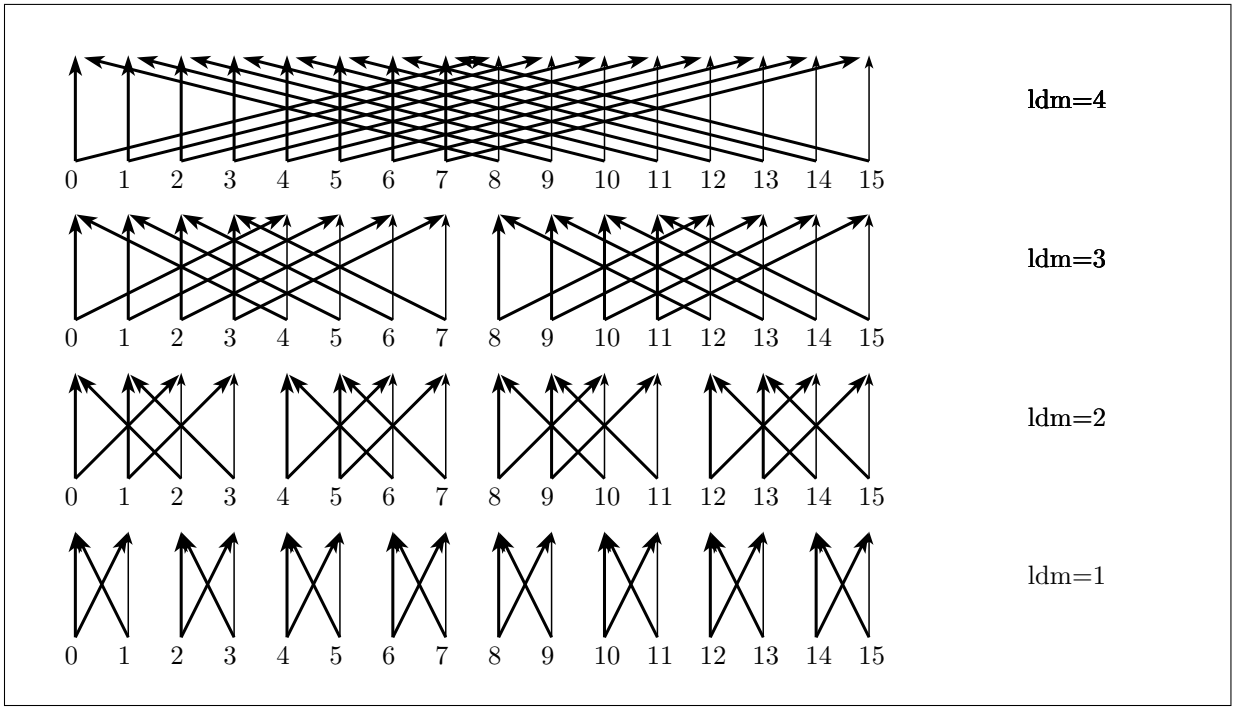

Figure 22.1-B: Data flow for the length-16, radix-2, decimation in time (DIT) transform. The stages are from bottom to top. Thin lines indicate a factor of minus one.

speedup by a factor of about *eight*. For smaller lengths the ratio approaches one.

The data flow diagram (*butterfly diagram*) for the radix-2 decimation in time (DIT) algorithm is shown in figure 22.1-B. The figure was created with the program [FXT: fft/butterfly-texpic-demo.cc]. The diagram for the decimation in frequency (DIF) algorithm is obtained by reversing the order of the steps. In the code, only the outermost loop has to be changed [FXT: `walsh_wak_dif2()` in walsh/walshwak2.h]:

```
template <typename Type>
void walsh_wak_dif2(Type *f, ulong ldn)
{
    const ulong n = (1UL<<ldn);
    for (ulong ldm=ldn; ldm>=1; --ldm)
    {
        const ulong m = (1UL<<ldm);
        const ulong mh = (m>>1);
        for (ulong r=0; r<n; r+=m)
        {
            ulong t1 = r;
            ulong t2 = r+mh;
            for (ulong j=0;  j<mh;  ++j, ++t1, ++t2)
            {
                Type u = f[t1];
                Type v = f[t2];
                f[t1] = u + v;
                f[t2] = u - v;
            }
        }
    }
}
```

A function that computes the k-th base function of the transform is [FXT: `walsh_wak_basefunc()` in walsh/walshbasefunc.h]:

```
template <typename Type>
void walsh_wak_basefunc(Type *f, ulong n, ulong k)
{
    for (ulong i=0; i<n; ++i)
    {
```

```
 6           ulong x = i & k;
 7           x = parity(x);
 8           f[i] = ( 0==x ? +1 : -1 );
 9       }
10   }
```

The basis functions are shown in figure 22.1-A. Note that the lowest row is (the signed version of) the *Thue-Morse sequence*, see section 1.16.1 on page 42.

Multi-dimensional Walsh transform

If one applies the row-column algorithm (see section 20.10.2 on page 423) to compute a two-dimensional $n \times m$ Walsh transform then the result is exactly the same as with a 1-dimensional $n \cdot m$ transform. That is, algorithmically nothing needs to be done for multidimensional Walsh transforms: a k-dimensional $n_1 \times n_2 \times \ldots \times n_k$-transform is identical to a 1-dimensional $n_1 \cdot n_2 \cdot \ldots \cdot n_k$-transform. The length-$2^n$ Walsh transform is identical to a n-dimensional length-2 Fourier transform.

22.2 Eigenvectors of the Walsh transform *

```
 0:  [ +5 +1 +1 +1 +1 +1 +1 +1 +1 +1 +1 +1 +1 +1 +1 +1 ]
 1:  [ +1 +3 +1 -1 +1 -1 +1 -1 +1 -1 +1 -1 +1 -1 +1 -1 ]
 2:  [ +1 +1 +3 -1 +1 +1 -1 -1 +1 +1 -1 -1 +1 +1 -1 -1 ]
 3:  [ +1 -1 -1 +5 +1 -1 -1 +1 +1 -1 -1 +1 +1 -1 -1 +1 ]
 4:  [ +1 +1 +1 +1 +3 -1 -1 -1 +1 +1 +1 +1 -1 -1 -1 -1 ]
 5:  [ +1 -1 +1 -1 -1 +5 -1 +1 +1 -1 +1 -1 -1 +1 -1 +1 ]
 6:  [ +1 +1 -1 -1 -1 -1 +5 +1 +1 +1 -1 -1 -1 -1 +1 +1 ]
 7:  [ +1 -1 -1 +1 -1 +1 +1 +3 +1 -1 -1 +1 -1 +1 +1 -1 ]
 8:  [ +1 +1 +1 +1 +1 +1 +1 +1 -5 -1 -1 -1 -1 -1 -1 -1 ]
 9:  [ +1 -1 +1 -1 +1 -1 +1 -1 -1 -3 -1 +1 -1 +1 -1 +1 ]
10:  [ +1 +1 -1 -1 +1 +1 -1 -1 -1 -1 -3 +1 -1 -1 +1 +1 ]
11:  [ +1 -1 -1 +1 +1 -1 -1 +1 -1 +1 +1 -5 -1 +1 +1 -1 ]
12:  [ +1 +1 +1 +1 -1 -1 -1 -1 -1 -1 -1 -1 -3 +1 +1 +1 ]
13:  [ +1 -1 +1 -1 -1 +1 -1 +1 -1 +1 -1 +1 +1 -5 +1 -1 ]
14:  [ +1 +1 -1 -1 -1 -1 +1 +1 -1 -1 +1 +1 +1 +1 -5 -1 ]
15:  [ +1 -1 -1 +1 -1 +1 +1 -1 -1 +1 +1 -1 +1 -1 -1 -3 ]
```

Figure 22.2-A: Eigenvectors of the length-16 Walsh transform (Walsh-Kronecker basis) as row vectors. The eigenvalues are $+1$ for the vectors $0 \ldots 7$ and -1 for the vectors $8 \ldots 16$. Linear combinations of vectors with the same eigenvalue e are again eigenvectors with eigenvalue e.

The Walsh transforms are self-inverse, so their eigenvalues can only be plus or minus one. Let a be a sequence and let $W(a)$ denote the Walsh transform of a. Set

$$u_+ \quad := \quad W(a) + a \tag{22.2-1}$$

Then

$$W(u_+) \quad = \quad W(W(a)) + W(a) \;=\; a + W(a) \;=\; +1 \cdot u_+ \tag{22.2-2}$$

That is, u_+ is an eigenvector of W with eigenvalue $+1$. Equivalently, $u_- := W(a) - a$ is an eigenvector with eigenvalue -1. Thereby, two eigenvectors can be obtained from an arbitrary nonzero sequence.

We are interested in a simple routine that for a Walsh transform of length n gives a set of n eigenvectors that span the n-dimensional space. With a routine that computes the k-th basis function of the transform we can obtain an eigenvector efficiently by simply adding a delta peak at position k to the basis function. The delta peak has to be scaled according to whether a positive or negative eigenvalue is desired and according to the normalization of the transform.

A suitable routine for the Walsh-Kronecker basis (whose basis functions are given in figure 22.1-A on page 448) is

```
1   void
2   walsh_wak_eigen(double *v, ulong ldn, ulong k)
3   // Eigenvectors of the Walsh transform (walsh_wak).
4   // Eigenvalues are +1 if k<n/2, else -1
5   {
6       ulong n = 1UL << ldn;
7       walsh_wak_basefunc(v, n, k);
8       double d = sqrt(n);
9       v[k] += (k<n/2 ? +d : -d);
10  }
```

This routine is given in [FXT: walsh/walsheigen.cc]. Figure 22.2-A was created with the program [FXT: fft/walsh-eigenvec-demo.cc].

Note that with the unnormalized transforms the eigenvalues are $\pm\sqrt{n}$.

22.3 The Kronecker product

The length-2 Walsh transform is equivalent to the multiplication of a 2-component vector by the matrix

$$\mathbf{W}_2 \quad = \quad \begin{bmatrix} +1 & +1 \\ +1 & -1 \end{bmatrix} \tag{22.3-1}$$

The length-4 Walsh transform corresponds to

$$\mathbf{W}_4 \quad = \quad \begin{bmatrix} +1 & +1 & +1 & +1 \\ +1 & -1 & +1 & -1 \\ +1 & +1 & -1 & -1 \\ +1 & -1 & -1 & +1 \end{bmatrix} \tag{22.3-2}$$

One might be tempted to write

$$\mathbf{W}_4 \quad = \quad \begin{bmatrix} +\mathbf{W}_2 & +\mathbf{W}_2 \\ +\mathbf{W}_2 & -\mathbf{W}_2 \end{bmatrix} \tag{22.3-3}$$

This idea can indeed be turned into a well-defined notation which is quite powerful when dealing with orthogonal transforms and their fast algorithms. Let $\mathbf{A}$ be an $m \times n$ matrix

$$\mathbf{A} \quad = \quad \begin{bmatrix} a_{0,0} & a_{0,1} & \cdots & a_{0,n-1} \\ a_{1,0} & a_{1,1} & \cdots & a_{1,n-1} \\ \vdots & \vdots & & \vdots \\ a_{m-1,0} & a_{m-1,1} & \cdots & a_{m-1,n-1} \end{bmatrix} \tag{22.3-4}$$

then the (right) *Kronecker product* (or *tensor product*) with a matrix $\mathbf{B}$ is

$$\mathbf{A} \otimes \mathbf{B} \quad := \quad \begin{bmatrix} a_{0,0}\mathbf{B} & a_{0,1}\mathbf{B} & \cdots & a_{0,n-1}\mathbf{B} \\ a_{1,0}\mathbf{B} & a_{1,1}\mathbf{B} & \cdots & a_{1,n-1}\mathbf{B} \\ \vdots & \vdots & & \vdots \\ a_{m-1,0}\mathbf{B} & a_{m-1,1}\mathbf{B} & \cdots & a_{m-1,n-1}\mathbf{B} \end{bmatrix} \tag{22.3-5}$$

There is no restriction on the dimensions of $\mathbf{B}$. If $\mathbf{B}$ is a $r \times s$ matrix then the dimensions of the given Kronecker product is $(mr) \times (ns)$, and $c_{k+ir,l+js} = a_{i,j} b_{k,l}$. The Kronecker product is not commutative, that is, $\mathbf{A} \otimes \mathbf{B} \neq \mathbf{B} \otimes \mathbf{A}$ in general.

For a scalar factor α the following relations are immediate:

$$(\alpha\mathbf{A}) \otimes \mathbf{B} = \alpha(\mathbf{A} \otimes \mathbf{B}) \quad (22.3\text{-}6a)$$

$$\mathbf{A} \otimes (\alpha\mathbf{B}) = \alpha(\mathbf{A} \otimes \mathbf{B}) \quad (22.3\text{-}6b)$$

The next relations are the same as for the ordinary matrix product. Distributivity (the matrices on both sides of a plus sign must be of the same dimensions):

$$(\mathbf{A} + \mathbf{B}) \otimes \mathbf{C} = \mathbf{A} \otimes \mathbf{C} + \mathbf{B} \otimes \mathbf{C} \quad (22.3\text{-}7a)$$

$$\mathbf{A} \otimes (\mathbf{B} + \mathbf{C}) = \mathbf{A} \otimes \mathbf{B} + \mathbf{A} \otimes \mathbf{C} \quad (22.3\text{-}7b)$$

Associativity:

$$\mathbf{A} \otimes (\mathbf{B} \otimes \mathbf{C}) = (\mathbf{A} \otimes \mathbf{B}) \otimes \mathbf{C} \quad (22.3\text{-}8)$$

The matrix product (indicated by a dot) of Kronecker products can be rewritten as

$$(\mathbf{A} \otimes \mathbf{B}) \cdot (\mathbf{C} \otimes \mathbf{D}) = (\mathbf{A} \cdot \mathbf{C}) \otimes (\mathbf{B} \cdot \mathbf{D}) \quad (22.3\text{-}9a)$$

$$(\mathbf{L}_1 \otimes \mathbf{R}_1) \cdot (\mathbf{L}_2 \otimes \mathbf{R}_2) \cdot \ldots \cdot (\mathbf{L}_n \otimes \mathbf{R}_n) = (\mathbf{L}_1 \cdot \mathbf{L}_2 \cdot \ldots \cdot \mathbf{L}_n) \otimes (\mathbf{R}_1 \cdot \mathbf{R}_2 \cdot \ldots \cdot \mathbf{R}_n) \quad (22.3\text{-}9b)$$

Set $\mathbf{L}_1 = \mathbf{L}_2 = \ldots = \mathbf{L}_n =: \mathbf{L}$ and $\mathbf{R}_1 = \mathbf{R}_2 = \ldots = \mathbf{R}_n =: \mathbf{R}$ in the latter relation to obtain

$$(\mathbf{L} \otimes \mathbf{R})^n = \mathbf{L}^n \otimes \mathbf{R}^n \quad (22.3\text{-}9c)$$

The Kronecker product of matrix products can be rewritten as

$$(\mathbf{A} \cdot \mathbf{B}) \otimes (\mathbf{C} \cdot \mathbf{D}) = (\mathbf{A} \otimes \mathbf{C}) \cdot (\mathbf{B} \otimes \mathbf{D}) \quad (22.3\text{-}10a)$$

$$(\mathbf{L}_1 \cdot \mathbf{R}_1) \otimes (\mathbf{L}_2 \cdot \mathbf{R}_2) \otimes \ldots \otimes (\mathbf{L}_n \cdot \mathbf{R}_n) = (\mathbf{L}_1 \otimes \mathbf{L}_2 \otimes \ldots \otimes \mathbf{L}_n) \cdot (\mathbf{R}_1 \otimes \mathbf{R}_2 \otimes \ldots \otimes \mathbf{R}_n) \quad (22.3\text{-}10b)$$

Here the matrices left and right from a dot must be compatible for ordinary matrix multiplication.

One has

$$(\mathbf{A} \otimes \mathbf{B})^T = \mathbf{A}^T \otimes \mathbf{B}^T \quad (22.3\text{-}11a)$$

$$(\mathbf{A} \otimes \mathbf{B})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{B}^{-1} \quad (22.3\text{-}11b)$$

If $\mathbf{A}$ and $\mathbf{B}$ are respectively $m \times n$ and $r \times s$ matrices then

$$\mathbf{A} \otimes \mathbf{B} = (\mathbf{I}_m \otimes \mathbf{B}) \cdot (\mathbf{A} \otimes \mathbf{I}_s) \quad (22.3\text{-}12a)$$

$$= (\mathbf{A} \otimes \mathbf{I}_r) \cdot (\mathbf{I}_n \otimes \mathbf{B}) \quad (22.3\text{-}12b)$$

where $\mathbf{I}_n$ is the $n \times n$ identity matrix. If $\mathbf{A}$ is $n \times n$ and $\mathbf{B}$ is $t \times t$ then

$$\det(\mathbf{A} \otimes \mathbf{B}) = \det(\mathbf{A})^t \det(\mathbf{B})^n \quad (22.3\text{-}13)$$

Back to the Walsh transform, we have $\mathbf{W}_1 = [1]$ and for $n = 2^k$, $n > 1$:

$$\mathbf{W}_n = \begin{bmatrix} +\mathbf{W}_{n/2} & +\mathbf{W}_{n/2} \\ +\mathbf{W}_{n/2} & -\mathbf{W}_{n/2} \end{bmatrix} = \mathbf{W}_2 \otimes \mathbf{W}_{n/2} \quad (22.3\text{-}14)$$

In order to see that this relation is the statement of a fast algorithm split the (to be transformed) vector x into halves

$$x = \begin{bmatrix} x_0 \\ x_1 \end{bmatrix} \quad (22.3\text{-}15)$$

and write out the matrix-vector product

$$\mathbf{W}_n\, x \quad = \quad \begin{bmatrix} \mathbf{W}_{n/2}\, x_0 + \mathbf{W}_{n/2}\, x_1 \\ \mathbf{W}_{n/2}\, x_0 - \mathbf{W}_{n/2}\, x_1 \end{bmatrix} \quad = \quad \begin{bmatrix} \mathbf{W}_{n/2}\, (x_0 + x_1) \\ \mathbf{W}_{n/2}\, (x_0 - x_1) \end{bmatrix} \tag{22.3-16}$$

That is, a length-n transform can be computed by two length-$n/2$ transforms of the sum and difference of the first and second half of x.

We define a notation equivalent to the product sign,

$$\bigotimes_{k=1}^{n} \mathbf{M}_k \quad := \quad \mathbf{M}_1 \otimes \mathbf{M}_2 \otimes \mathbf{M}_3 \otimes \ldots \otimes \mathbf{M}_n \tag{22.3-17}$$

where the empty product equals a 1×1 matrix with entry 1. When $\mathbf{A} = \mathbf{B}$ in relation 22.3-11b we have $(\mathbf{A} \otimes \mathbf{A})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{A}^{-1}$, $(\mathbf{A} \otimes \mathbf{A} \otimes \mathbf{A})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{A}^{-1} \otimes \mathbf{A}^{-1}$ and so on. That is,

$$\left(\bigotimes_{k=1}^{n} \mathbf{A} \right)^{-1} \quad = \quad \bigotimes_{k=1}^{n} \mathbf{A}^{-1} \tag{22.3-18}$$

For the Walsh transform we have

$$\mathbf{W}_n \quad = \quad \bigotimes_{k=1}^{\log_2(n)} \mathbf{W}_2 \tag{22.3-19}$$

and

$$\mathbf{W}_n^{-1} \quad = \quad \bigotimes_{k=1}^{\log_2(n)} \mathbf{W}_2^{-1} \tag{22.3-20}$$

The latter relation isn't that exciting as $\mathbf{W}_2^{-1} = \mathbf{W}_2$ for the Walsh transform. However, it also holds when the inverse transform is different from the forward transform. Thereby, given a fast algorithm for some transform in form of a Kronecker product, the fast algorithm for the backward transform is immediate.

The *direct sum* of two matrices is defined as

$$\mathbf{A} \oplus \mathbf{B} \quad := \quad \begin{bmatrix} \mathbf{A} & 0 \\ 0 & \mathbf{B} \end{bmatrix} \tag{22.3-21}$$

In general $\mathbf{A} \oplus \mathbf{B} \neq \mathbf{B} \oplus \mathbf{A}$. As an analogue to the sum sign we have

$$\bigoplus_{k=1}^{n} \mathbf{A} \quad := \quad \mathbf{I}_n \otimes \mathbf{A} \tag{22.3-22}$$

where $\mathbf{I}_n$ is the $n \times n$ identity matrix. The matrix $\mathbf{I}_n \otimes \mathbf{A}$ consists of n copies of $\mathbf{A}$ that lie on the diagonal. The Kronecker product can be used to derive properties of unitary transforms, see [221].

22.4 A variant of the Walsh transform *

All operations necessary for the Walsh transform are cheap: loads, stores, additions and subtractions. The memory access pattern is a major concern with direct mapped cache, as we have verified comparing the first two implementations in this chapter. Even the one found to be superior due to its more localized access is guaranteed to have a performance problem as soon as the array is long enough: all accesses are separated by a power-of-two distance and cache misses will occur beyond a certain limit. Rather bizarre attempts like inserting 'pad data' have been reported in order to mitigate the problem. The Gray code permutation described in section 2.8 on page 105 suggests an interesting solution where the sub-arrays are always accessed in mutually reversed order [FXT: walsh/walshgray.h]:

```
template <typename Type>
void walsh_gray(Type *f, ulong ldn)
// Gray variant of the Walsh transform.
// Radix-2 decimation in frequency (DIF) algorithm
{
    const ulong n = (1UL<<ldn);

    for (ulong ldm=ldn; ldm>0; --ldm) // dif
    {
        const ulong m = (1UL<<ldm);
        for (ulong r=0; r<n; r+=m)
        {
            ulong t1 = r;
            ulong t2 = r + m - 1;
            for (  ;  t1<t2;  ++t1,--t2)
            {
                Type u = f[t1];
                Type v = f[t2];
                f[t1] = u + v;
                f[t2] = u - v;
            }
        }
    }
}
```

The transform is not self-inverse, however, the inverse transform can be implemented easily be reversing the steps:

```
template <typename Type>
void inverse_walsh_gray(Type *f, ulong ldn)
// Inverse of walsh_gray().
// Radix-2 decimation in time (DIT) algorithm.
{
    const ulong n = (1UL<<ldn);

    for (ulong ldm=1; ldm<=ldn; ++ldm) // dit
    {
        const ulong m = (1UL<<ldm);
        for (ulong r=0; r<n; r+=m)
        {
            ulong t1 = r;
            ulong t2 = r + m - 1;
            for (  ;  t1<t2;  ++t1,--t2)
            {
                Type u = f[t1];
                Type v = f[t2];
                f[t1] = u + v;
                f[t2] = u - v;
            }
        }
    }
}
```

Using Q for the `grs_negate()` routine (described below), W_k for `walsh_wak()`, W_g for `walsh_gray()` and G for `gray_permute()` then

$$W_k \;=\; Q\,W_g\,G^{-1} \;=\; G\,W_g^{-1}\,Q \tag{22.4-1}$$

That is, either of the following two sequences of statements

```
    { inverse_gray_permute(f, n);  walsh_gray(f, ldn);  grs_negate(f, n); }
    { grs_negate(f, n);  inverse_walsh_gray(f, ldn);  gray_permute(f, n); }
```

is equivalent to the call `walsh_wak(f, ldn)`. The function [FXT: `grs_negate()` in aux1/grsnegate.h] changes signs for certain elements:

```
template <typename Type>
void grs_negate(Type *f, ulong n)
// Negate elements at indices where the Golay-Rudin-Shapiro is negative.
{
    for (ulong k=0; k<n; ++k)
    {
        if ( grs_negative_q(k) )  f[k] = -f[k];
    }
```

```
9   }
```

The function `grs_negative_q()` is described in section 1.16.5 on page 45.

It turns out that the Gray-variant only wins on machines where the memory clock speed is significantly lower than the CPU. While a call to `walsh_gray()` alone is never slower than the `walsh_wak()` routine the additional steps often cause too much overhead.

22.5 Higher radix Walsh transforms

A generator for short-length Walsh (wak) transforms is given as [FXT: fft/gen-walsh-demo.cc]. It can create code for DIF and DIT transforms. For example, the code for the 4-point DIF transforms is

```
1   template <typename Type>
2   inline void
3   short_walsh_wak_dif_4(Type *f)
4   {
5       Type t0, t1, t2, t3;
6       t0 = f[0];
7       t1 = f[1];
8       t2 = f[2];
9       t3 = f[3];
10      sumdiff( t0, t2 );
11      sumdiff( t1, t3 );
12      sumdiff( t0, t1 );
13      sumdiff( t2, t3 );
14      f[0] = t0;
15      f[1] = t1;
16      f[2] = t2;
17      f[3] = t3;
18  }
```

To make the code more readable we use function [FXT: `sumdiff()` in aux0/sumdiff.h]:

```
1   template <typename Type>
2   static inline void sumdiff(Type &a, Type &b)
3   // {a, b}  <--| {a+b, a-b}
4   { Type t=a-b; a+=b; b=t; }
```

We further need a variant that transforms elements which are not contiguous but lie apart by a distance s:

```
1   template <typename Type>
2   inline void
3   short_walsh_wak_dif_4(Type *f, ulong s)
4   {
5       Type t0, t1, t2, t3;
6       {
7       ulong x = 0;
8       t0 = f[x];  x += s;
9       t1 = f[x];  x += s;
10      t2 = f[x];  x += s;
11      t3 = f[x];
12      }
13      sumdiff( t0, t2 );
14      sumdiff( t1, t3 );
15      sumdiff( t0, t1 );
16      sumdiff( t2, t3 );
17      {
18      ulong x = 0;
19      f[x] = t0;  x += s;
20      f[x] = t1;  x += s;
21      f[x] = t2;  x += s;
22      f[x] = t3;
23      }
24  }
```

The short DIF transforms are given in [FXT: walsh/shortwalshwakdif.h], DIT variants in [FXT: walsh/shortwalshwakdit.h]. A radix-4 DIF transform using these ingredients is [FXT: walsh/walshwak4.h]:

```
1   template <typename Type>
```

```
void walsh_wak_dif4(Type *f, ulong ldn)
// Transform wrt. to Walsh-Kronecker basis (wak-functions).
// Radix-4 decimation in frequency (DIF) algorithm.
// Self-inverse.
{
    const ulong n = (1UL<<ldn);

    if ( n<=2 )
    {
        if ( n==2 )  short_walsh_wak_dif_2(f);
        return;
    }

    for (ulong ldm=ldn; ldm>3; ldm-=2)
    {
        ulong m = (1UL<<ldm);
        ulong m4 = (m>>2);
        for (ulong r=0;  r<n;  r+=m)
        {
            for (ulong j=0;  j<m4;  j++)  short_walsh_wak_dif_4(f+j+r, m4);
        }
    }

    if ( ldn & 1 )  // n is not a power of 4, need a radix-8 step
    {
        for (ulong i0=0; i0<n; i0+=8)  short_walsh_wak_dif_8(f+i0);
    }
    else
    {
        for (ulong i0=0; i0<n; i0+=4)  short_walsh_wak_dif_4(f+i0);
    }
}
```

With the implementation radix-8 DIF transform some care must be taken to choose the correct final step size [FXT: walsh/walshwak8.h]:

```
template <typename Type>
void walsh_wak_dif8(Type *f, ulong ldn)
// Transform wrt. to Walsh-Kronecker basis (wak-functions).
// Radix-8 decimation in frequency (DIF) algorithm.
// Self-inverse.
{
    const ulong n = (1UL<<ldn);

    if ( n<=4 )
    {
        switch (n )
        {
        case 4:  short_walsh_wak_dif_4(f);  break;
        case 2:  short_walsh_wak_dif_2(f);  break;
        }
        return;
    }

    const ulong xx = 4;
    ulong ldm;
    for (ldm=ldn; ldm>xx; ldm-=3)
    {
        ulong m = (1UL<<ldm);
        ulong m8 = (m>>3);
        for (ulong r=0;  r<n;  r+=m)
        {
            for (ulong j=0;  j<m8;  j++)  short_walsh_wak_dif_8(f+j+r, m8);
        }
    }

    switch ( ldm )
    {
    case 4:
        for (ulong i0=0; i0<n; i0+=16)  short_walsh_wak_dif_16(f+i0);
        break;
    case 3:
        for (ulong i0=0; i0<n; i0+=8)  short_walsh_wak_dif_8(f+i0);
        break;
    case 2:
        for (ulong i0=0; i0<n; i0+=4)  short_walsh_wak_dif_4(f+i0);
        break;
    }
```

```
43  }
```

Performance

For the performance comparison we include a matrix variant of the Walsh transform [FXT: walsh/walshwakmatrix.h]:

```
1   template <typename Type>
2   void walsh_wak_matrix(Type *f, ulong ldn)
3   {
4       ulong ldc = (ldn>>1);
5       ulong ldr = ldn-ldc;  // ldr>=ldc
6       ulong nc = (1UL<<ldc);
7       ulong nr = (1UL<<ldr);  // nrow >= ncol
8
9       for (ulong r=0; r<nr; ++r)  walsh_wak_dif4(f+r*nc, ldc);
10      transpose2(f, nr, nc);
11      for (ulong c=0; c<nc; ++c)  walsh_wak_dif4(f+c*nr, ldr);
12      transpose2(f, nc, nr);
13  }
```

The transposition routine is given in [FXT: aux2/transpose2.h]. We only use even powers of two so the transposition is that of a square matrix.

As for dyadic convolutions we do not need the data in a particular order so we also include a version of the matrix algorithm that omits the final transposition:

```
1   template <typename Type>
2   void walsh_wak_matrix_1(Type *f, ulong ldn, int is)
3   {
4       ulong ldc = (ldn>>1);
5       ulong ldr = ldn-ldc;  // ldr>=ldc
6       if ( is<0 )  swap2(ldr, ldc);  // inverse
7       ulong nc = (1UL<<ldc);
8       ulong nr = (1UL<<ldr);  // nrow >= ncol
9
10      for (ulong r=0; r<nr; ++r)  walsh_wak_dif4(f+r*nc, ldc);
11      transpose2(f, nr, nc);
12      for (ulong c=0; c<nc; ++c)  walsh_wak_dif4(f+c*nr, ldr);
13  }
```

The following calls give (up to normalization) the mutually inverse transforms:

```
1       walsh_wak_matrix_1(f, ldn, +1);
2       walsh_wak_matrix_1(f, ldn, -1);
```

We do not consider the range of transform lengths $n < 128$ where unrolled routines and the radix-4 algorithm consistently win. Figure 22.5-A shows a comparison of the routines given so far. There are clearly two regions to distinguish: firstly, the region were the transforms fit into the first-level data cache (which is 64 kilobyte, corresponding to `ldn` = 13). Secondly, the region where `ldn` > 13 and the performance becomes more and more memory bound.

In the first region the radix-4 routine is the fastest. The radix-8 routine comes close but, somewhat surprisingly, never wins.

In the second region the matrix version is the best. However, for very large sizes its performance could be better. Note that with odd `ldn` (not shown) its performance drops significantly due to the more expensive transposition operation. The transposition is clearly the bottleneck. One can use (machine specific) optimizations for the transposition to further improve the performance.

In the next section we give an algorithm that avoids the transposition completely and consistently outperforms the matrix algorithm.

```
 8 == ldn;   MemSize ==     2 kB  ==  256 doubles;   rep == 976563
              walsh_wak_dif2(f,ldn);    dt=    2.49551    MB/s= 6114   rel=          1
              walsh_wak_dif4(f,ldn);    dt=    1.56806    MB/s= 9731   rel=   0.628352 *
              walsh_wak_dif8(f,ldn);    dt=    1.57419    MB/s= 9693   rel=    0.63081
            walsh_wak_matrix(f,ldn);    dt=    2.28047    MB/s= 6691   rel=    0.91383
      walsh_wak_matrix_1(f,ldn,+1);     dt=    1.94357    MB/s= 7851   rel=   0.778827
                 walsh_gray(f,ldn);     dt=    2.39791    MB/s= 6363   rel=   0.960888

10 == ldn;   MemSize ==     8 kB  ==  1024 doubles;   rep == 195313
              walsh_wak_dif2(f,ldn);    dt=    2.26683    MB/s= 6731   rel=          1
              walsh_wak_dif4(f,ldn);    dt=    1.47338    MB/s=10356   rel=   0.649977 *
              walsh_wak_dif8(f,ldn);    dt=    1.65262    MB/s= 9233   rel=   0.729044
            walsh_wak_matrix(f,ldn);    dt=    1.91859    MB/s= 7953   rel=   0.846378
      walsh_wak_matrix_1(f,ldn,+1);     dt=    1.69215    MB/s= 9017   rel=   0.746485
                 walsh_gray(f,ldn);     dt=    2.18454    MB/s= 6985   rel=     0.9637

12 == ldn;   MemSize ==    32 kB  ==  4096 doubles;   rep == 20345
              walsh_wak_dif2(f,ldn);    dt=     1.0884    MB/s= 7010   rel=          1
              walsh_wak_dif4(f,ldn);    dt=   0.723136    MB/s=10550   rel=   0.664403 *
              walsh_wak_dif8(f,ldn);    dt=   0.790313    MB/s= 9654   rel=   0.726124
            walsh_wak_matrix(f,ldn);    dt=    1.01233    MB/s= 7536   rel=   0.930112
      walsh_wak_matrix_1(f,ldn,+1);     dt=   0.926387    MB/s= 8236   rel=   0.851146
                 walsh_gray(f,ldn);     dt=    1.05364    MB/s= 7241   rel=   0.968062

14 == ldn;   MemSize ==   128 kB  ==  16384 doubles;   rep == 2180
              walsh_wak_dif2(f,ldn);    dt=    1.17042    MB/s= 3260   rel=          1
              walsh_wak_dif4(f,ldn);    dt=    1.14861    MB/s= 3321   rel=   0.981368
              walsh_wak_dif8(f,ldn);    dt=    1.08501    MB/s= 3516   rel=   0.927026
            walsh_wak_matrix(f,ldn);    dt=   0.669182    MB/s= 5701   rel=   0.571747
      walsh_wak_matrix_1(f,ldn,+1);     dt=   0.552063    MB/s= 6910   rel=   0.471681 *
                 walsh_gray(f,ldn);     dt=    1.00794    MB/s= 3785   rel=    0.86118

16 == ldn;   MemSize ==   512 kB  ==  65536 doubles;   rep == 477
              walsh_wak_dif2(f,ldn);    dt=    1.40004    MB/s= 2726   rel=          1
              walsh_wak_dif4(f,ldn);    dt=    1.70347    MB/s= 2240   rel=    1.21673
              walsh_wak_dif8(f,ldn);    dt=    1.12997    MB/s= 3377   rel=   0.807095
            walsh_wak_matrix(f,ldn);    dt=   0.801902    MB/s= 4759   rel=   0.572769
      walsh_wak_matrix_1(f,ldn,+1);     dt=   0.628073    MB/s= 6076   rel=   0.448609 *
                 walsh_gray(f,ldn);     dt=    1.18693    MB/s= 3215   rel=   0.847779

18 == ldn;   MemSize ==     2 MB  ==  256 K doubles;   rep == 106
              walsh_wak_dif2(f,ldn);    dt=    2.61599    MB/s= 1459   rel=          1
              walsh_wak_dif4(f,ldn);    dt=    2.55153    MB/s= 1496   rel=   0.975359
              walsh_wak_dif8(f,ldn);    dt=     1.9791    MB/s= 1928   rel=   0.756538
            walsh_wak_matrix(f,ldn);    dt=    1.77306    MB/s= 2152   rel=   0.677776
      walsh_wak_matrix_1(f,ldn,+1);     dt=    1.14735    MB/s= 3326   rel=   0.438591 *
                 walsh_gray(f,ldn);     dt=    2.51539    MB/s= 1517   rel=   0.961545

20 == ldn;   MemSize ==     8 MB  ==  1024 K doubles;   rep == 24
              walsh_wak_dif2(f,ldn);    dt=    2.64158    MB/s= 1454   rel=          1
              walsh_wak_dif4(f,ldn);    dt=     2.8532    MB/s= 1346   rel=    1.08011
              walsh_wak_dif8(f,ldn);    dt=    2.34867    MB/s= 1635   rel=   0.889113
            walsh_wak_matrix(f,ldn);    dt=    1.88431    MB/s= 2038   rel=   0.713327
      walsh_wak_matrix_1(f,ldn,+1);     dt=    1.21084    MB/s= 3171   rel=   0.458376 *
                 walsh_gray(f,ldn);     dt=    2.58747    MB/s= 1484   rel=   0.979514

22 == ldn;   MemSize ==    32 MB  ==  4096 K doubles;   rep == 5
              walsh_wak_dif2(f,ldn);    dt=    2.43537    MB/s= 1445   rel=          1
              walsh_wak_dif4(f,ldn);    dt=    2.82337    MB/s= 1247   rel=    1.15932
              walsh_wak_dif8(f,ldn);    dt=    2.07422    MB/s= 1697   rel=   0.851708
            walsh_wak_matrix(f,ldn);    dt=    1.99251    MB/s= 1767   rel=   0.818155
      walsh_wak_matrix_1(f,ldn,+1);     dt=    1.22719    MB/s= 2868   rel=   0.503901 *
                 walsh_gray(f,ldn);     dt=    2.38611    MB/s= 1475   rel=   0.979776

24 == ldn;   MemSize ==   128 MB  ==  16384 K doubles;   rep == 1
              walsh_wak_dif2(f,ldn);    dt=    2.10939    MB/s= 1456   rel=          1
              walsh_wak_dif4(f,ldn);    dt=    2.61517    MB/s= 1175   rel=    1.23977
              walsh_wak_dif8(f,ldn);    dt=    2.11508    MB/s= 1452   rel=     1.0027
            walsh_wak_matrix(f,ldn);    dt=    2.16597    MB/s= 1418   rel=    1.02683
      walsh_wak_matrix_1(f,ldn,+1);     dt=    1.28349    MB/s= 2393   rel=   0.608466 *
                 walsh_gray(f,ldn);     dt=    2.04744    MB/s= 1500   rel=   0.970635
```

Figure 22.5-A: Relative speed of different implementations of the Walsh (wak) transform. The transforms were run '`rep`' times for each measurement, the quantity '`dt`' gives the elapsed time for `rep` transforms of the given type. The quantity '`MB/s`' gives the memory transfer rate as if a radix-2 algorithm was used, it equals '`Memsize`' times '`ldn`' divided by the time elapsed for a single transform. The '`rel`' gives the performance relative to the radix-2 version, smaller values mean better performance.

22.6 Localized Walsh transforms

A decimation in time (DIT) algorithm combines the two halves of the array, then the halves of the two halves, the halves of each quarter, and so on. With each step the whole array is accessed which leads to the performance drop as soon as the array does not fit into the cache.

22.6.1 The method of localization

One can reorganize the algorithm as follows: combine the two halves of the array and postpone further processing of the upper half, then combine the halves of the lowers half and again postpone processing of its upper half. Repeat until size two is reached. Then use the algorithm at the postponed parts, starting with the smallest (last postponed).

The scheme can be sketched for size 16, as follows:

```
hhhhhhhhhhhhhhhh
hhhhhhhh44444444
hhhh333344444444
hh22333344444444
```

The letters 'h' denote places processed before any recursive call. The blocks of twos, threes and fours denote postponed blocks. The Walsh transform is thereby decomposed into a sequence of Haar transforms (see figure 23.6-A on page 494). The algorithm described is most easily implemented via recursion:

```
1  template <typename Type>
2  void walsh_wak_loc_dit2(Type *f, ulong ldn)
3  {
4      if ( ldn<1 )  return;
5
6      // Recursion:
7      for (ulong ldm=1; ldm<ldn; ++ldm)  walsh_wak_loc_dit2(f+(1UL<<ldm), ldm);
8
9      for (ulong ldm=1; ldm<=ldn; ++ldm)
10     {
11         const ulong m = (1UL<<ldm);
12         const ulong mh = (m>>1);
13         for (ulong t1=0, t2=mh;  t1<mh;  ++t1, ++t2)  sumdiff(f[t1], f[t2]);
14     }
15 }
```

22.6.2 Optimizing the routine

Optimizations are obtained by avoiding recursions for small sizes. We use a radix-4 algorithm as soon as the transform size is smaller or equal to the cache size and we avoid recursion for tiny transforms [FXT: walsh/walshwakloc2.h]:

```
1  template <typename Type>
2  void walsh_wak_loc_dit2(Type *f, ulong ldn)
3  {
4      if ( ldn<=13 )  // parameter: (2**13)*sizeof(Type) <= L1-cache
5      {
6          walsh_wak_dif4(f,ldn);  // note: DIF version, result is the same
7          return;
8      }
9
10     // Recursion:
11     short_walsh_wak_dit_2(f+2);  // ldm==1
12     short_walsh_wak_dit_4(f+4);  // ldm==2
13     short_walsh_wak_dit_8(f+8);  // ldm==3
14     short_walsh_wak_dit_16(f+16);  // ldm==4
15     for (ulong ldm=5; ldm<ldn; ++ldm)  walsh_wak_loc_dit2(f+(1UL<<ldm), ldm);
16
17     for (ulong ldm=1; ldm<=ldn; ++ldm)
18     {
19         const ulong m = (1UL<<ldm);
```

```
20          const ulong mh = (m>>1);
21          for (ulong t1=0, t2=mh;  t1<mh;  ++t1, ++t2)  sumdiff(f[t1], f[t2]);
22      }
23  }
```

A decimation in frequency (DIF) version is obtained by reversion of the steps:

```
1   template <typename Type>
2   void walsh_wak_loc_dif2(Type *f, ulong ldn)
3   {
4       if ( ldn<=13 )  // parameter: (2**13)*sizeof(Type) <= L1-cache
5       {
6           walsh_wak_dif4(f,ldn);
7           return;
8       }
9
10      for (ulong ldm=ldn; ldm>=1; --ldm)
11      {
12          const ulong m = (1UL<<ldm);
13          const ulong mh = (m>>1);
14          for (ulong t1=0, t2=mh;  t1<mh;  ++t1, ++t2)
15          {
16              Type u = f[t1];
17              Type v = f[t2];
18              f[t1] = u + v;
19              f[t2] = u - v;
20          }
21      }
22
23      // Recursion:
24      short_walsh_wak_dif_2(f+2);  // ldm==1
25      short_walsh_wak_dif_4(f+4);  // ldm==2
26      short_walsh_wak_dif_8(f+8);  // ldm==3
27      short_walsh_wak_dif_16(f+16);  // ldm==4
28      for (ulong ldm=5; ldm<ldn; ++ldm)  walsh_wak_loc_dif2(f+(1UL<<ldm), ldm);
29  }
```

The double loop in the algorithm is a reversed Haar transform, see chapter 23 on page 483. The double loop in the DIF algorithm is a transposed reversed Haar transform. The (generated) short-length transforms are given in the files [FXT: walsh/shortwalshwakdif.h] and [FXT: walsh/shortwalshwakdit.h]. For example, the length-8, decimation in frequency routine is

```
1   template <typename Type>
2   inline void
3   short_walsh_wak_dif_8(Type *f)
4   {
5    Type t0, t1, t2, t3, t4, t5, t6, t7;
6    t0 = f[0]; t1 = f[1]; t2 = f[2]; t3 = f[3];
7    t4 = f[4]; t5 = f[5]; t6 = f[6]; t7 = f[7];
8    sumdiff( t0, t4 ); sumdiff( t1, t5 ); sumdiff( t2, t6 ); sumdiff( t3, t7 );
9    sumdiff( t0, t2 ); sumdiff( t1, t3 ); sumdiff( t4, t6 ); sumdiff( t5, t7 );
10   sumdiff( t0, t1 ); sumdiff( t2, t3 ); sumdiff( t4, t5 ); sumdiff( t6, t7 );
11   f[0] = t0; f[1] = t1; f[2] = t2; f[3] = t3;
12   f[4] = t4; f[5] = t5; f[6] = t6; f[7] = t7;
13  }
```

The used strategy leads to a very favorable memory access pattern that results in excellent performance for large transforms. Figure 22.6-A shows a comparison between the localized transforms and the matrix algorithm. Small sizes are omitted because the localized algorithm has the very same speed as the radix-4 algorithm it falls back to. The localized algorithms are the clear winners, even against the matrix algorithm with only one transposition. For very large transforms the decimation in frequency version is very slightly faster. This is due to the fact that it starts with smaller chunks of data so more of the data is in cache when the larger sub-arrays are accessed.

The localized algorithm can easily be implemented for transforms where a radix-2 step is known. Section 24.9 on page 515 gives the fast Hartley transform variant of the localized algorithm.

One can develop similar routines with higher radix. However, a radix-4 version was found to be slower than the given routines. A certain speedup can be achieved by unrolling and prefetching. We use the C-type `double` whose size is 8 bytes. Substitute the double loop in the DIF version (that is, the Haar transform) by

```
   14 == ldn;   MemSize ==   128 kB  ==  16384 doubles;   rep == 2180
              walsh_wak_matrix(f,ldn);    dt=   0.672327    MB/s= 5674   rel=          1
         walsh_wak_matrix_1(f,ldn,+1);    dt=   0.555851    MB/s= 6863   rel=   0.826756
            walsh_wak_loc_dit2(f,ldn);    dt=   0.498558    MB/s= 7652   rel=   0.741541 *
            walsh_wak_loc_dif2(f,ldn);    dt=   0.533746    MB/s= 7148   rel=   0.793878

   16 == ldn;   MemSize ==   512 kB  ==  65536 doubles;   rep == 477
              walsh_wak_matrix(f,ldn);    dt=   0.919579    MB/s= 4150   rel=          1
         walsh_wak_matrix_1(f,ldn,+1);    dt=   0.692488    MB/s= 5511   rel=   0.753049
            walsh_wak_loc_dit2(f,ldn);    dt=   0.653256    MB/s= 5842   rel=   0.710386 *
            walsh_wak_loc_dif2(f,ldn);    dt=   0.670104    MB/s= 5695   rel=   0.728707

   18 == ldn;   MemSize ==     2 MB  ==  256 K doubles;   rep == 106
              walsh_wak_matrix(f,ldn);    dt=     2.2111    MB/s= 1726   rel=          1
         walsh_wak_matrix_1(f,ldn,+1);    dt=    1.36827    MB/s= 2789   rel=   0.618819
            walsh_wak_loc_dit2(f,ldn);    dt=   0.938006    MB/s= 4068   rel=   0.424225
            walsh_wak_loc_dif2(f,ldn);    dt=   0.927804    MB/s= 4113   rel=   0.419611 *

   20 == ldn;   MemSize ==     8 MB  ==  1024 K doubles;   rep == 24
              walsh_wak_matrix(f,ldn);    dt=    2.31178    MB/s= 1661   rel=          1
         walsh_wak_matrix_1(f,ldn,+1);    dt=    1.42614    MB/s= 2693   rel=   0.616901
            walsh_wak_loc_dit2(f,ldn);    dt=    1.11847    MB/s= 3433   rel=   0.483811
            walsh_wak_loc_dif2(f,ldn);    dt=    1.11142    MB/s= 3455   rel=   0.480765 *

   22 == ldn;   MemSize ==    32 MB  ==  4096 K doubles;   rep == 5
              walsh_wak_matrix(f,ldn);    dt=    2.00573    MB/s= 1755   rel=          1
         walsh_wak_matrix_1(f,ldn,+1);    dt=    1.23695    MB/s= 2846   rel=   0.616707
            walsh_wak_loc_dit2(f,ldn);    dt=    1.16461    MB/s= 3022   rel=   0.580644
            walsh_wak_loc_dif2(f,ldn);    dt=    1.16164    MB/s= 3030   rel=   0.579162 *

   24 == ldn;   MemSize ==   128 MB  ==  16384 K doubles;   rep == 1
              walsh_wak_matrix(f,ldn);    dt=    2.16536    MB/s= 1419   rel=          1
         walsh_wak_matrix_1(f,ldn,+1);    dt=    1.28455    MB/s= 2392   rel=   0.593226
            walsh_wak_loc_dit2(f,ldn);    dt=    1.10769    MB/s= 2773   rel=   0.511552
            walsh_wak_loc_dif2(f,ldn);    dt=    1.10601    MB/s= 2778   rel=   0.510775 *
```

Figure 22.6-A: Speed comparison between localized and matrix algorithms for the Walsh transform.

```
1   // machine specific prefetch instruction:
2   #define PREF(p,o)  asm volatile ("prefetchw " #o "(%0) " : : "r" (p) )
3
4       ulong ldm;
5       for (ldm=ldn; ldm>=6; --ldm)
6       {
7           const ulong m = (1UL<<ldm);
8           const ulong mh = (m>>1);
9           PREF(f, 0);  PREF(f+mh, 0);
10          PREF(f, 64);  PREF(f+mh, 64);
11          PREF(f, 128);  PREF(f+mh, 128);
12          PREF(f, 192);  PREF(f+mh, 192);
13
14          for (ulong t1=0, t2=mh;  t1<mh;  t1+=8, t2+=8)
15          {
16              double *p1 = f + t1,  *p2 = f + t2;
17              PREF(p1, 256);  PREF(p2, 256);
18
19              double u0 = f[t1+0], v0 = f[t2+0];
20              double u1 = f[t1+1], v1 = f[t2+1];
21              double u2 = f[t1+2], v2 = f[t2+2];
22              double u3 = f[t1+3], v3 = f[t2+3];
23              sumdiff(u0, v0); f[t1+0] = u0; f[t2+0] = v0;
24              sumdiff(u1, v1); f[t1+1] = u1; f[t2+1] = v1;
25              sumdiff(u2, v2); f[t1+2] = u2; f[t2+2] = v2;
26              sumdiff(u3, v3); f[t1+3] = u3; f[t2+3] = v3;
27
28              double u4 = f[t1+4], v4 = f[t2+4];
29              double u5 = f[t1+5], v5 = f[t2+5];
30              double u6 = f[t1+6], v6 = f[t2+6];
31              double u7 = f[t1+7], v7 = f[t2+7];
32              sumdiff(u4, v4); f[t1+4] = u4; f[t2+4] = v4;
33              sumdiff(u5, v5); f[t1+5] = u5; f[t2+5] = v5;
34              sumdiff(u6, v6); f[t1+6] = u6; f[t2+6] = v6;
35              sumdiff(u7, v7); f[t1+7] = u7; f[t2+7] = v7;
36          }
37      }
38
39      for (  ; ldm>=1; --ldm)
```

```
40      {
41          const ulong m = (1UL<<ldm);
42          const ulong mh = (m>>1);
43          for (ulong t1=0, t2=mh;  t1<mh;  ++t1, ++t2)  sumdiff(f[t1], f[t2]);
44      }
```

The following list gives the speed ratio between the optimized and the unoptimized DIF routine:

```
14 == ldn;   MemSize ==   128 kB;     ratio = 1.24252
16 == ldn;   MemSize ==   512 kB;     ratio = 1.43568
18 == ldn;   MemSize ==     2 MB;     ratio = 1.23875
20 == ldn;   MemSize ==     8 MB;     ratio = 1.21012
22 == ldn;   MemSize ==    32 MB;     ratio = 1.19939
24 == ldn;   MemSize ==   128 MB;     ratio = 1.18245
```

For sizes that are out of (level-2) cache most of the speedup is due to the memory prefetch.

22.6.3 Iterative versions of the algorithms

```
 DIF        DIF            DIT        DIT
start      length         start      length
......     1.....         ....1.     ....1.
....1.     ....1.         ...11.     ...11.
...1..     ...1..         ...1..     ...1..
...11.     ....1.         ..1.1.     ....1.
..1...     ..1...         ..111.     ...11.
..1.1.     ....1.         ..11..     ..1...
..11..     ...1..         ..1...     ..1...
..111.     ....1.         .1..1.     ....1.
.1....     .1....         .1.11.     ...11.
.1..1.     ....1.         .1.1..     ...1..
.1.1..     ...1..         .11.1.     ....1.
.1.11.     ....1.         .1111.     ...11.
.11...     ..1...         .111..     ..1...
.11.1.     ....1.         .11...     .1....
.111..     ...1..         .1....     1.....
.1111.     ....1.         ......
```

Figure 22.6-B: Binary values of the start index and length of the Haar transforms in the iterative version of the localized DIF (left) and DIT (right) transform. Dots are used for zeros.

In the DIF algorithm the Haar transforms are executed at positions $f+2$, $f+4$, $f+6$, ... and the length of the transform at position $f+s$ is determined by the lowest set bit in s. Additionally, a full-length Haar transform has to be done at the beginning. As C++ code:

```
1   template <typename Type>
2   inline void haar_dif2(Type *f, ulong n)
3   {
4       for (ulong m=n; m>=2; m>>=1)
5       {
6           const ulong mh = (m>>1);
7           for (ulong t1=0, t2=mh;  t1<mh;  ++t1, ++t2)  sumdiff(f[t1], f[t2]);
8       }
9   }
10
11  template <typename Type>
12  void loc_dif2(Type *f, ulong n)
13  {
14      haar_dif2(f, n);
15      for (ulong z=2; z<n; z+=2)  haar_dif2(f+z,  (z&-z));
16  }
```

Note that the routines now take the length of the transform as second argument, not its base-2 logarithm.

With the DIT algorithm matters are slightly more complicated. A pattern can be observed by printing the binary expansions of the starting position and length of the transforms shown in figure 22.6-B (created with [FXT: fft/locrec-demo.cc]). The lengths are again determined by the lowest bit of the start position. And we have also seen the pattern in the left column: the reversed binary words in reversed (subset-) lexicographic order, see figure 1.28-A on page 71. The implementation is quite concise:

```
1   template <typename Type>
2   inline void haar_dit2(Type *f, ulong n)
```

```
3   {
4       for (ulong m=1; m<=n; m<<=1)
5       {
6           const ulong mh = (m>>1);
7           for (ulong t1=0, t2=mh;  t1<mh;  ++t1, ++t2)  sumdiff(f[t1], f[t2]);
8       }
9   }
10
11  template <typename Type>
12  void loc_dit2(Type f, ulong n)
13  {
14      for (ulong z=2, u=1; z<n; z+=2)
15      {
16          ulong s = u<<1;
17          haar_dit2(f+s,  (s&-s));
18          u = prev_lexrev(u);
19      }
20      haar_dit2(f, n);
21  }
```

The routines are slightly slower than the recursive version because they do not fall back to the full Walsh transforms when the transform size is small.

The DIT scheme is a somewhat surprising application of the seemingly esoteric routine [FXT: `prev_lexrev()` in bits/bitlex.h]. Plus we have found a recursive algorithm for the generation of the binary words in (subset-) lexicographic order [FXT: bits/bitlex-rec-demo.cc]:

```
1   void bitlex_b(ulong f, ulong n)
2   {
3       for (ulong m=1; m<n; m<<=1)  bitlex_b(f+m, m);
4       print_bin("  ", f, ldn);
5   }
```

22.7 Dyadic (XOR) convolution

Dyadic convolution has XOR where the usual one has plus

The *dyadic convolution* of the sequences a and b is the sequence h defined by

$$h_\tau := \sum_{i \oplus j = \tau} a_i \, b_j \tag{22.7-1a}$$

$$= \sum_{i} a_i \, b_{i \oplus \tau} \tag{22.7-1b}$$

where the symbol '$\oplus$' stands for bit-wise XOR operator. All three sequence must be of the same length that is a power of 2. The dyadic convolution could rightfully be called *XOR-convolution*.

The semi-symbolic scheme of the convolution is shown in figure 22.7-A. The table is equivalent to the one (for cyclic convolution) given in figure 21.1-A on page 428. The dyadic convolution can be used for the multiplication of hypercomplex numbers as shown in section 37.16 on page 810.

A fast algorithm for the computation of the dyadic convolution uses the Walsh transform [FXT: `dyadic_convolution()` in walsh/dyadiccnvl.h]:

```
1   template <typename Type>
2   void dyadic_convolution(Type * restrict f, Type * restrict g, ulong ldn)
3   // Dyadic convolution (XOR-convolution):  h[] of f[] and g[]:
4   //   h[k] = sum( i XOR j == k,  f[i]*g[k] )
5   // Result is written to g[].
6   // ldn := base-2 logarithm of the array length
7   {
8       walsh_wak(f, ldn);
9       walsh_wak(g, ldn);
10      const ulong n = (1UL<<ldn);
11      for (ulong k=0; k<n; ++k)  g[k] *= f[k];
12      walsh_wak(g, ldn);
13  }
```

```
          +--  0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
          |
          0:   0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
          1:   1  0  3  2   5  4  7  6   9  8 11 10  13 12 15 14
          2:   2  3  0  1   6  7  4  5  10 11  8  9  14 15 12 13
          3:   3  2  1  0   7  6  5  4  11 10  9  8  15 14 13 12

          4:   4  5  6  7   0  1  2  3  12 13 14 15   8  9 10 11
          5:   5  4  7  6   1  0  3  2  13 12 15 14   9  8 11 10
          6:   6  7  4  5   2  3  0  1  14 15 12 13  10 11  8  9
          7:   7  6  5  4   3  2  1  0  15 14 13 12  11 10  9  8

          8:   8  9 10 11  12 13 14 15   0  1  2  3   4  5  6  7
          9:   9  8 11 10  13 12 15 14   1  0  3  2   5  4  7  6
         10:  10 11  8  9  14 15 12 13   2  3  0  1   6  7  4  5
         11:  11 10  9  8  15 14 13 12   3  2  1  0   7  6  5  4

         12:  12 13 14 15   8  9 10 11   4  5  6  7   0  1  2  3
         13:  13 12 15 14   9  8 11 10   5  4  7  6   1  0  3  2
         14:  14 15 12 13  10 11  8  9   6  7  4  5   2  3  0  1
         15:  15 14 13 12  11 10  9  8   7  6  5  4   3  2  1  0
```

Figure 22.7-A: Semi-symbolic scheme for the dyadic convolution of two length-16 sequences.

```
          +--  0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
          |
          0:   0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
          1:   1  0  3  2   5  4  7  6   9  8 11 10  13 12 15 14
          2:   2  3  0  1   6  7  4  5  10 11  8  9  14 15 12 13
          3:   3  2  1  0   7  6  5  4  11 10  9  8  15 14 13 12

          4:   4  5  6  7   0  1  2  3  12 13 14 15   8  9 10 11
          5:   5  4  7  6   1  0  3  2  13 12 15 14   9  8 11 10
          6:   6  7  4  5   2  3  0  1  14 15 12 13  10 11  8  9
          7:   7  6  5  4   3  2  1  0  15 14 13 12  11 10  9  8

          8:   8  9 10 11  12 13 14 15   0- 1- 2- 3-  4- 5- 6- 7-
          9:   9  8 11 10  13 12 15 14   1- 0- 3- 2-  5- 4- 7- 6-
         10:  10 11  8  9  14 15 12 13   2- 3- 0- 1-  6- 7- 4- 5-
         11:  11 10  9  8  15 14 13 12   3- 2- 1- 0-  7- 6- 5- 4-

         12:  12 13 14 15   8  9 10 11   4- 5- 6- 7-  0- 1- 2- 3-
         13:  13 12 15 14   9  8 11 10   5- 4- 7- 6-  1- 0- 3- 2-
         14:  14 15 12 13  10 11  8  9   6- 7- 4- 5-  2- 3- 0- 1-
         15:  15 14 13 12  11 10  9  8   7- 6- 5- 4-  3- 2- 1- 0-
```

Figure 22.7-B: Semi-symbolic scheme for the dyadic equivalent of the negacyclic convolution. Negative contributions to a bucket have a minus appended.

An scheme similar to that of the negacyclic convolution is shown in figure 22.7-B. It can be obtained via

```
1       walsh_wal_dif2_core(f, ldn);
2       walsh_wal_dif2_core(g, ldn);
3       ulong n = (1UL<<ldn);
4       for (ulong i=0,j=n-1; i<j; --j,++i)  fht_mul(f[i], f[j], g[i], g[j], 0.5);
5       walsh_wal_dit2_core(g, ldn);
```

where `fht_mul()` is the operation used for the convolution with fast Hartley transforms [FXT: convolution/fhtmulsqr.h]:

```
1   template <typename Type>
2   static inline  void
3   fht_mul(Type xi, Type xj, Type &yi, Type &yj, double v)
4   // yi <-- v*( 2*xi*xj + xi*xi - xj*xj )
5   // yj <-- v*( 2*xi*xj - xi*xi + xj*xj )
6   {
7       Type h1p = xi,  h1m = xj;
8       Type s1 = h1p + h1m,  d1 = h1p - h1m;
9       Type h2p = yi,  h2m = yj;
10      yi = (h2p * s1 + h2m * d1) * v;
11      yj = (h2m * s1 - h2p * d1) * v;
12  }
```

22.8 The Walsh transform: Walsh-Paley basis

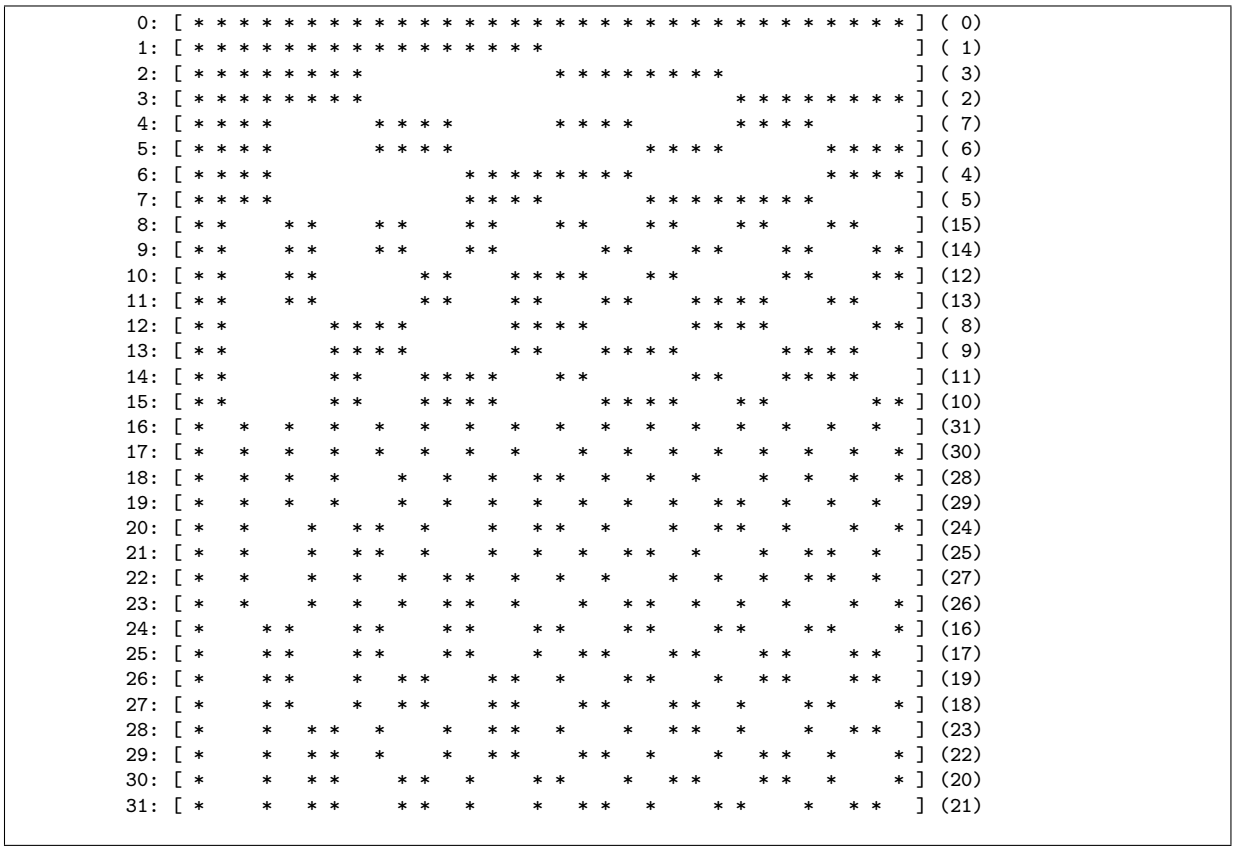

Figure 22.8-A: Basis functions for the Walsh transform (Walsh-Paley basis). Asterisks denote the value $+1$, blank entries denote -1.

A Walsh transform with a different (ordering of the) basis can be obtained by [FXT: walsh/walshpal.h]:

```
1   template <typename Type>
2   void walsh_pal(Type *f, ulong ldn)
3   {
4       const ulong n = 1UL<<ldn;
5       revbin_permute(f, n);
6       walsh_wak(f, ldn);
7   // =^=
8   //    walsh_wak(f, ldn);
9   //    revbin_permute(f, n);
10  }
```

The basis functions are shown in figure 22.8-A. Actually one can also apply the revbin permutation before the transform. That is,

$$W_p \;=\; W_k\, R \;=\; R\, W_k \tag{22.8-1}$$

One has for W_p

$$W_p \;=\; G\, W_p\, G \;=\; G^{-1}\, W_p\, G^{-1} \tag{22.8-2}$$

$$\;=\; Z\, W_p\, Z \;=\; Z^{-1}\, W_p\, Z^{-1} \tag{22.8-3}$$

where Z denotes the zip permutation (see section 2.5 on page 101) and G denotes the Gray permutation (see section 2.8 on page 105).

A function that computes the k-th base function of the transform is [FXT: `walsh_pal_basefunc()` in walsh/walshbasefunc.h]:

```
template <typename Type>
void walsh_pal_basefunc(Type *f, ulong n, ulong k)
{
    k = revbin(k, ld(n));
    for (ulong i=0; i<n; ++i)
    {
        ulong x = i & k;
        x = parity(x);
        f[i] = ( 0==x ? +1 : -1 );
    }
}
```

22.9 Sequency ordered Walsh transforms

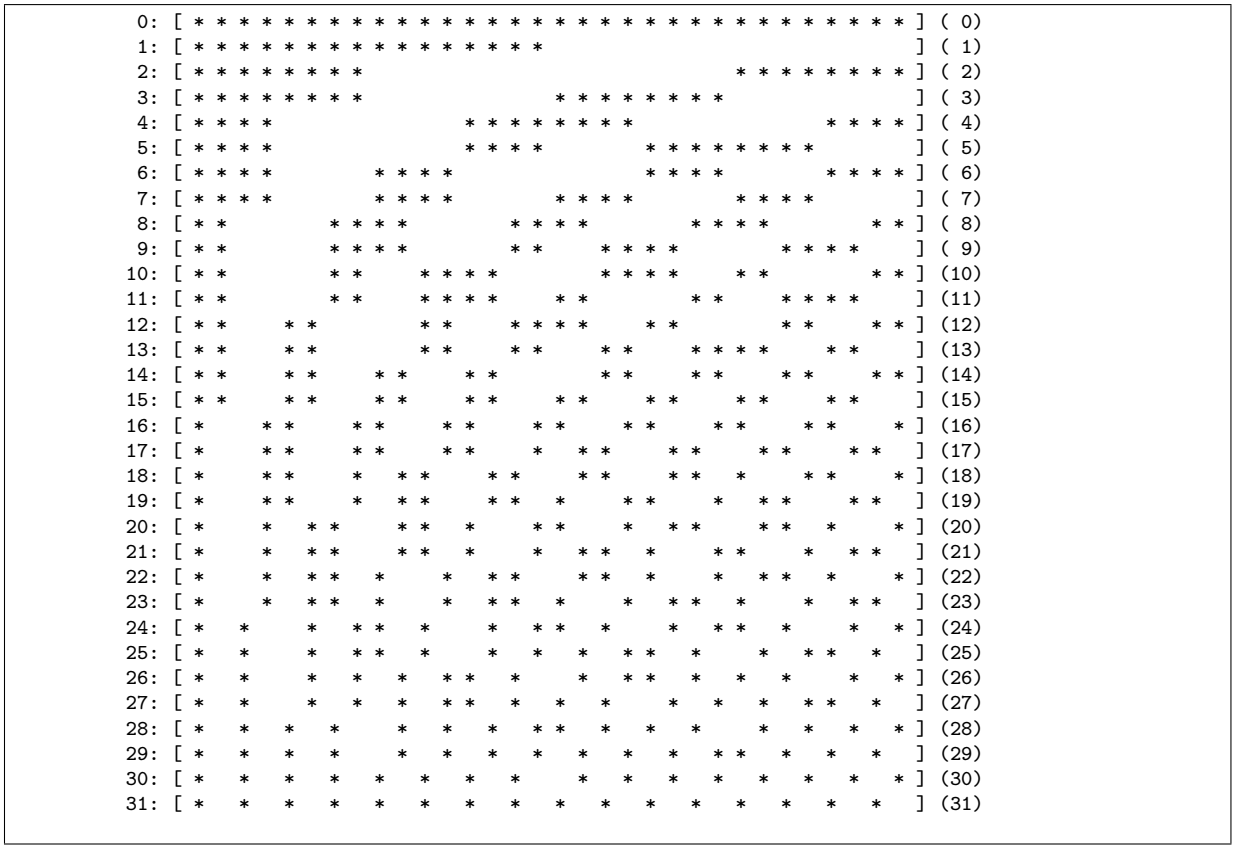

Figure 22.9-A: Basis functions for the sequency-ordered Walsh transform (Walsh-Kacmarz basis). Asterisks denote the value $+1$, blank entries denote -1.

The term corresponding to the frequency of the Fourier basis functions is the so-called *sequency* of the Walsh functions, the number of the changes of sign of the individual functions. Note that the sequency of a signal with frequency f usually is $2\,f$.

If the basis functions shall be ordered by their sequency one can use

```
    const ulong n = (1UL<<ldn);
    walsh_wak(f, ldn);
    revbin_permute(f, n);
    inverse_gray_permute(f, n);
```

That is

$$W_w \;=\; G^{-1}\,R\,W_k \;=\; W_k\,R\,G \tag{22.9-1}$$

A function that computes the k-th base function of the transform is [FXT: `walsh_wal_basefunc()` in walsh/walshbasefunc.h]:

```
template <typename Type>
void walsh_wal_basefunc(Type *f, ulong n, ulong k)
{
    k = revbin(k, ld(n)+1);
    k = gray_code(k);
//    // =^=
//    k = revbin(k, ld(n));
//    k = rev_gray_code(k);
    for (ulong i=0; i<n; ++i)
    {
        ulong x = i & k;
        x = parity(x);
        f[i] = ( 0==x ? +1 : -1 );
    }
}
```

A version of the transform that avoids the Gray permutation is based on [FXT: walsh/walshwal.h]

```
template <typename Type>
void walsh_wal_dif2_core(Type *f, ulong ldn)
// Core routine for sequency ordered Walsh transform.
// Radix-2 decimation in frequency (DIF) algorithm.
{
    const ulong n = (1UL<<ldn);
    for (ulong ldm=ldn; ldm>=2; --ldm)
    {
        const ulong m = (1UL<<ldm);
        const ulong mh = (m>>1);
        const ulong m4 = (mh>>1);
        for (ulong r=0; r<n; r+=m)
        {
            ulong j;
            for (j=0; j<m4; ++j)
            {
                ulong t1 = r+j;
                ulong t2 = t1+mh;
                double u = f[t1];
                double v = f[t2];
                f[t1] = u + v;
                f[t2] = u - v;
            }

            for (  ; j<mh; ++j)
            {
                ulong t1 = r+j;
                ulong t2 = t1+mh;
                double u = f[t1];
                double v = f[t2];
                f[t1] = u + v;
                f[t2] = v - u;  // reversed
            }
        }
    }

    if ( ldn )
    {
        // ulong ldm=1;
        const ulong m = 2; //(1UL<<ldm);
        const ulong mh = 1; //(m>>1);
        for (ulong r=0; r<n; r+=m)
        {
            ulong j = 0;
//            for (ulong j=0; j<mh; ++j)
            {
                ulong t1 = r+j;
                ulong t2 = t1+mh;
```

```
49                  double u = f[t1];
50                  double v = f[t2];
51                  f[t1] = u + v;
52                  f[t2] = u - v;
53              }
54          }
55      }
56  }
```

The transform still needs the revbin permutation:

```
1   template <typename Type>
2   inline void walsh_wal(Type *f, ulong ldn)
3   {
4       revbin_permute(f, (1UL<<ldn));
5       walsh_wal_dif2_core(f, ldn);
6   //  =^=
7   //    walsh_wal_dit2_core(f, ldn);
8   //    revbin_permute(f, (1UL<<ldn));
9   }
```

A decimation in time (DIT) version of the core-routine is also given in [FXT: walsh/walshwal.h]. The procedure `gray_permute()` is given in section 2.8 on page 105.

The sequence of statements

```
1       walsh_gray(f, ldn);  grs_negate(f, n);  revbin_permute(f, n);
```

is equivalent to `walsh_wal(f, ldn)` and might be faster for large arrays. We have

$$W_w \;=\; R\,Q\,W_g \;=\; W_g^{-1}\,R\,Q \tag{22.9-2}$$

22.9.1 Even/odd ordering of sequencies

An alternative ordering of the base functions (first even sequencies ascending then odd sequencies descending [FXT: `walsh_wal_rev()` in walsh/walshwalrev.h]) can be obtained by either of the following sequences of statements (with `n=1UL<<ldn`):

```
1       { revbin_permute(f, n);  gray_permute(f, n);  walsh_wak(f,ldn); }
2       { walsh_wak(f,ldn);  inverse_gray_permute(f, n); revbin_permute(f, n); }
3
4       { zip_rev(f, n);  walsh_wal(f, ldn); }
5       { walsh_wal(f, ldn);  unzip_rev(f, n); }
6
7       { walsh_wak(f, ldn);  inverse_gray_permute(f, n);  revbin_permute(f, n); }
8
9       { revbin_permute(f, n);  walsh_gray(f, ldn);  grs_negate(f, n); }
```

That is,

$$\begin{aligned}
\overline{W}_w \;&=\; W_k\,G\,R \;=\; R\,G^{-1}\,W_k && (22.9\text{-}3)\\
&=\; W_w\,\overline{Z} \;=\; \overline{Z}^{-1}\,W_w && (22.9\text{-}4)\\
&=\; Q\,W_g\,R && (22.9\text{-}5)
\end{aligned}$$

However, an implementation that is more efficient uses the core-routines that have the Gray permutation 'absorbed' [FXT: `walsh_wal_rev()` in walsh/walshwalrev.h]:

```
1   template <typename Type>
2   inline void walsh_wal_rev(Type *f, ulong ldn)
3   {
4       revbin_permute(f, (1UL<<ldn));
5       walsh_wal_dit2_core(f, ldn);
6       // =^=
7   //    walsh_wal_dif2_core(f, ldn);
8   //    revbin_permute(f, n);
9   }
```

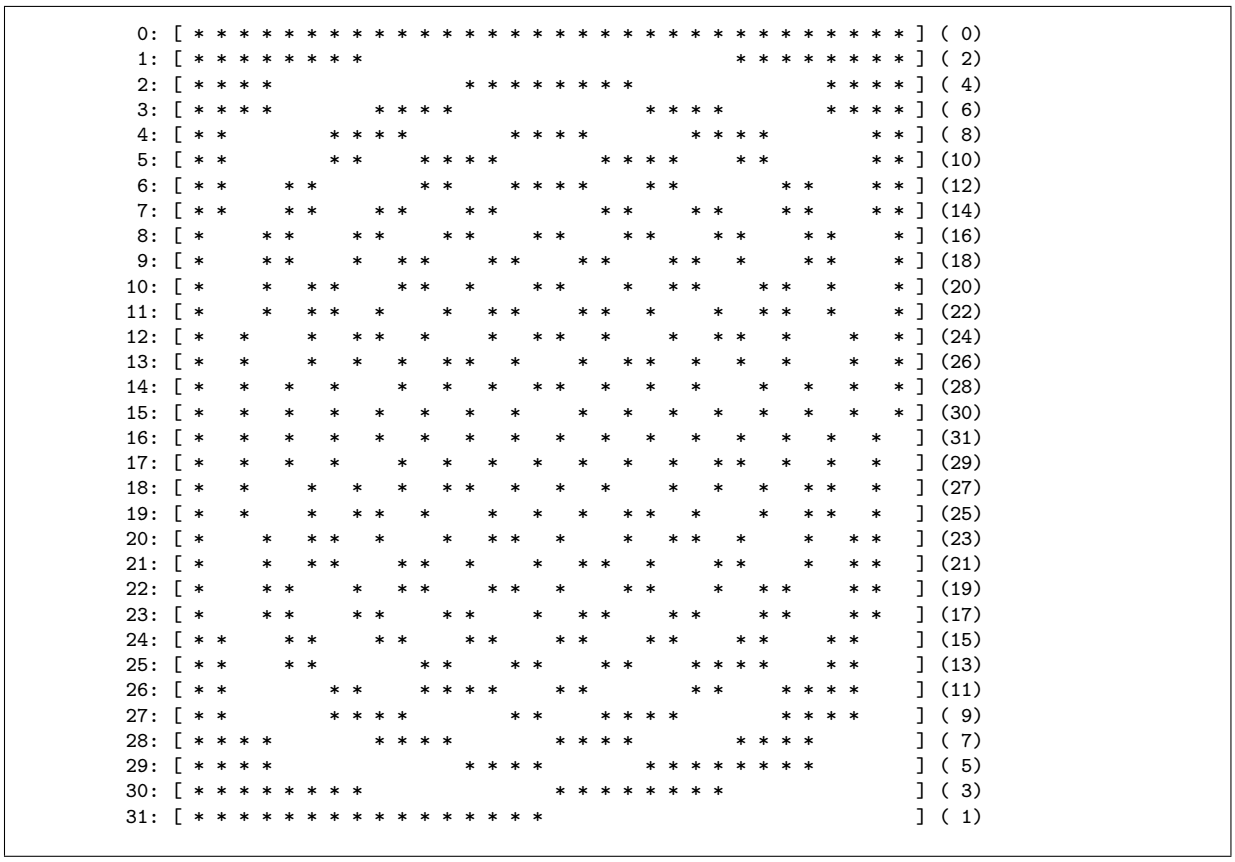

Figure 22.9-B: Basis functions for the reversed sequency ordered Walsh transform. Asterisks denote the value $+1$, blank entries denote -1.

This implementation uses the fact that

$$\overline{W}_w \quad = \quad R\, W_w\, R \tag{22.9-6}$$

The following sequence of statements computes the same transform:

```
{ revbin_permute(f, n);  walsh_gray(f, ldn);  grs_negate(f, n); }
```

Similar relations as for the transform with Walsh-Paley basis (22.8-2 and 22.8-3 on page 465) hold for W_w:

$$W_w \quad = \quad G\, W_w\, G \;=\; G^{-1}\, W_w\, G^{-1} \tag{22.9-7}$$

$$\quad = \quad \overline{Z}\, W_w\, \overline{Z} \;=\; \overline{Z}^{-1}\, W_w\, \overline{Z}^{-1} \tag{22.9-8}$$

The k-th base function of the transform can be computed as [FXT: `walsh_wal_rev_basefunc()` in walsh/walshbasefunc.h]

```
1   template <typename Type>
2   void walsh_wal_rev_basefunc(Type *f, ulong n, ulong k)
3   {
4       k = revbin(k, ld(n));
5       k = gray_code(k);
6       // =^=
7   //    k = rev_gray_code(k);
8   //    k = revbin(k, ld(n));
9       for (ulong i=0; i<n; ++i)
10      {
11          ulong x = i & k;
12          x = parity(x);
13          f[i] = ( 0==x ? +1 : -1 );
14      }
15  }
```

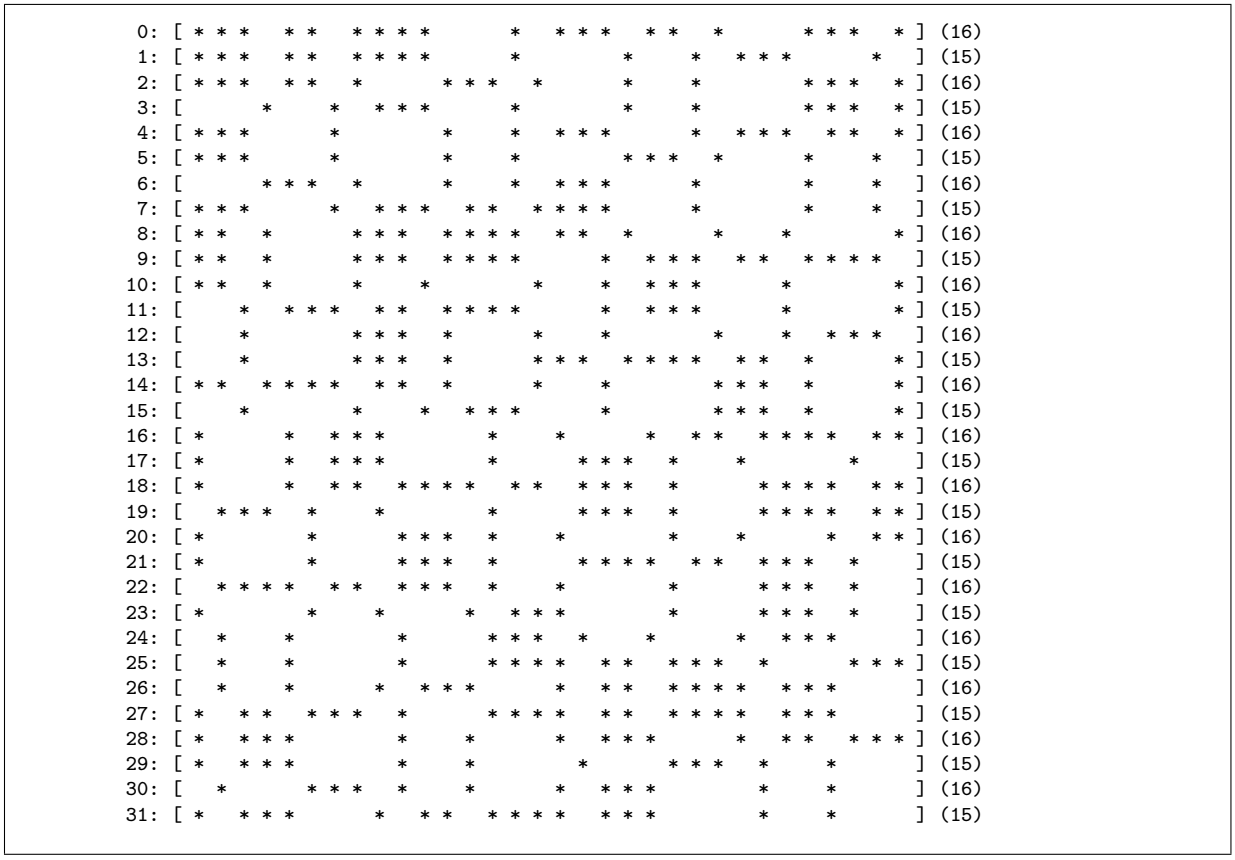

Figure 22.9-C: Basis functions for a self-inverse Walsh transform that has sequencies $n/2$ and $n/2-1$ only. Asterisks denote the value $+1$, blank entries denote -1.

22.9.2 Transforms with sequencies $n/2$ or $n/2-1$

The next variant of the Walsh transform has the interesting feature that the basis functions for a length-n transform have only sequencies $n/2$ and $n/2-1$ at the even and odd indices, respectively. The transform is self-inverse (the basis is shown in figure 22.9-C) and can be obtained via [FXT: `walsh_q1()` in walsh/walshq.h]

```
1   template <typename Type>
2   void walsh_q1(Type *f, ulong ldn)
3   {
4       ulong n = 1UL << ldn;
5       grs_negate(f, n);
6       walsh_gray(f, ldn);
7       revbin_permute(f, n);
8   }
```

A different transform with sequency $n/2$ for the first half of the basis, sequency $n/2-1$ for the second half ([FXT: `walsh_q2()` in walsh/walshq.h], basis shown in figure 22.9-D) is computed by

```
1   template <typename Type>
2   void walsh_q2(Type *f, ulong ldn)
3   {
4       ulong n = 1UL << ldn;
5       revbin_permute(f, n);
6       grs_negate(f, n);
7       walsh_gray(f, ldn);
8   // =^=
9   //    grs_negate(f, n);
10  //    revbin_permute(f, n);
11  //    walsh_gray(f, ldn);
```

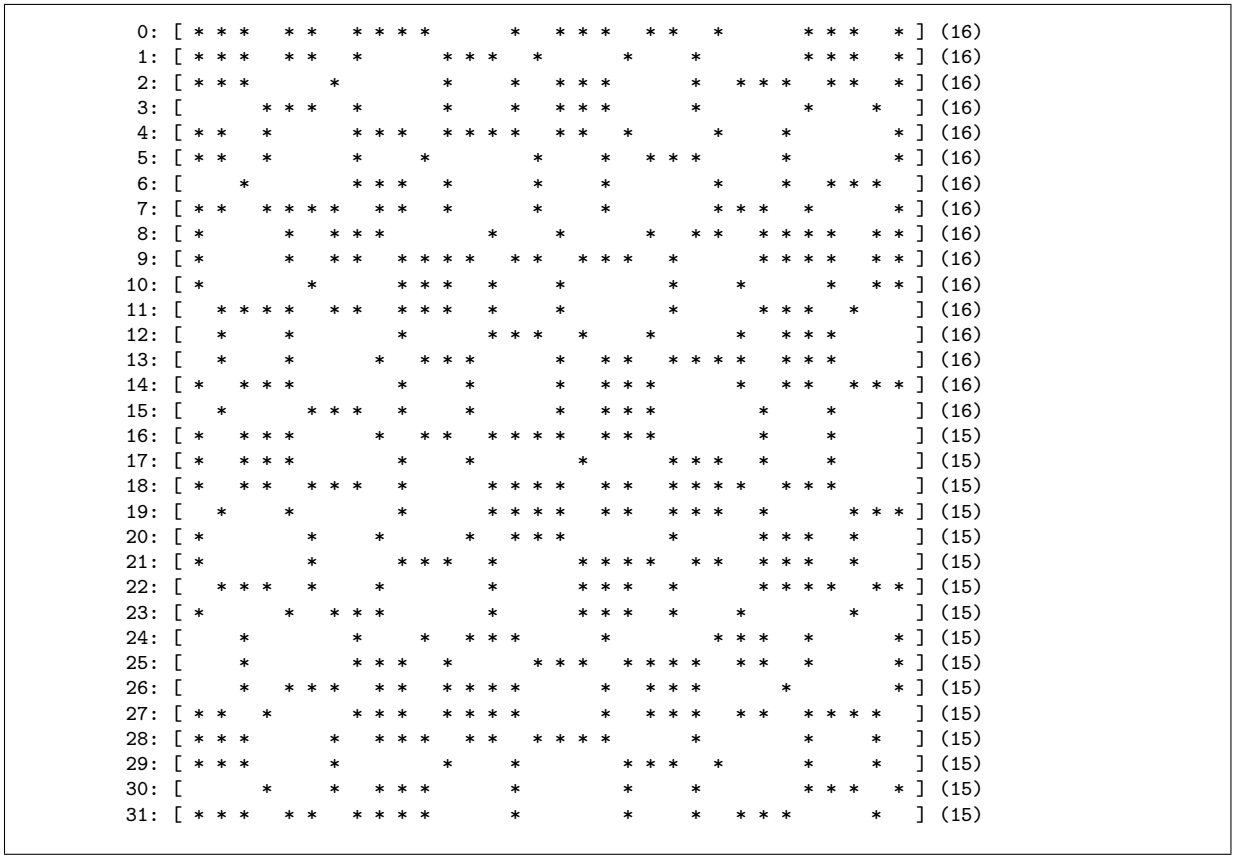

Figure 22.9-D: Basis functions for a self-inverse Walsh transform (second form) that has sequencies $n/2$ and $n/2-1$ only. Asterisks denote the value $+1$, blank entries denote -1.

```
12  }
```

One has:

$$W_{q2} \quad = \quad R\,W_{q1}\,R \tag{22.9-9}$$

The base functions of the transforms can be computed as [FXT: walsh/walshbasefunc.h]

```
1   template <typename Type>
2   void walsh_q1_basefunc(Type *f, ulong n, ulong k)
3   {
4       ulong qk = (grs_negative_q(k) ? 1 : 0);
5       k = gray_code(k);
6       k = revbin(k, ld(n));
7       for (ulong i=0; i<n; ++i)
8       {
9           ulong x = i & k;
10          x = parity(x);
11          ulong qi = (grs_negative_q(i) ? 1 : 0);
12          x ^= (qk ^ qi);
13          f[i] = ( 0==x ? +1 : -1 );
14      }
15  }
```

and

```
1   template <typename Type>
2   void walsh_q2_basefunc(Type *f, ulong n, ulong k)
3   {
4       ulong qk = (grs_negative_q(k) ? 1 : 0);
5       k = revbin(k, ld(n));
```

```
6       k = gray_code(k);
7       for (ulong i=0; i<n; ++i)
8       {
9           ulong x = i & k;
10          x = parity(x);
11          ulong qi = (grs_negative_q(i) ? 1 : 0);
12          x ^= (qk ^ qi);
13          f[i] = ( 0==x ? +1 : -1 );
14      }
15  }
```

The function `grs_negative_q()` is described in section 1.16.5 on page 45.

22.10 Slant transform

The *slant transform* can be implemented using a Walsh Transform and just a little pre/post-processing [FXT: walsh/slant.cc]:

```
1   void slant(double *f, ulong ldn)
2   {
3       walsh_wak(f, ldn);
4
5       ulong n = 1UL<<ldn;
6       for (ulong ldm=0; ldm<ldn-1; ++ldm)
7       {
8           ulong m = 1UL<<ldm;  // m = 1, 2, 4, 8, ..., n/4
9           double N = m*2, N2 = N*N;
10          double a = sqrt(3.0*N2/(4.0*N2-1.0));
11          double b = sqrt(1.0-a*a); // == sqrt((N2-1)/(4*N2-1));
12          for (ulong j=m; j<n-1; j+=4*m)
13          {
14              ulong t1 = j;
15              ulong t2 = j + m;
16              double f1 = f[t1],  f2 = f[t2];
17              f[t1] = a * f1 - b * f2;
18              f[t2] = b * f1 + a * f2;
19          }
20      }
21  }
```

Apart from the Walsh transform only an amount of work linear with the array size has to be done: the inner loop accesses the elements in strides of 4, 8, 16, ..., 2^{n-1}.

The inverse transform is:

```
1   void inverse_slant(double *f, ulong ldn)
2   {
3       ulong n = 1UL<<ldn;
4       ulong ldm=ldn-2;
5       do
6       {
7           ulong m = 1UL<<ldm;  // m = n/4, n/2, ..., 4, 2, 1
8           double N = m*2, N2 = N*N;
9           double a = sqrt(3.0*N2/(4.0*N2-1.0));
10          double b = sqrt(1.0-a*a); // == sqrt((N2-1)/(4*N2-1));
11          for (ulong j=m; j<n-1; j+=4*m)
12          {
13              ulong t1 = j;
14              ulong t2 = j + m;
15              double f1 = f[t1],  f2 = f[t2];
16              f[t1] = b * f2 + a * f1;
17              f[t2] = a * f2 - b * f1;
18          }
19      }
20      while ( ldm-- );
21
22      walsh_wak(f, ldn);
23  }
```

A sequency ordered version of the transform can be implemented as follows:

```
1   void slant_seq(double *f, ulong ldn)
```

```
2   {
3       slant(f, ldn);
4
5       ulong n = 1UL<<ldn;
6       inverse_gray_permute(f, n);
7       unzip_rev(f, n);
8       revbin_permute(f, n);
9   }
```

This implementation can be optimized by combining the involved permutations, see [264].

The inverse is obtained by calling the inverse operations in reversed order:

```
1   void inverse_slant_seq(double *f, ulong ldn)
2   {
3       ulong n = 1UL<<ldn;
4       revbin_permute(f, n);
5       zip_rev(f, n);
6       gray_permute(f, n);
7
8       inverse_slant(f, ldn);
9   }
```

22.11 Arithmetic transform

How to make a arithmetic transform out of a Walsh transform:
'*Forward: replace* `(u+v)` *by* `(u)`*, and* `(u-v)` *by* `(v-u)`*.*
Backward: replace `(u+v)` *by* `(u)`*, and* `(u-v)` *by* `(u+v)`*.* '

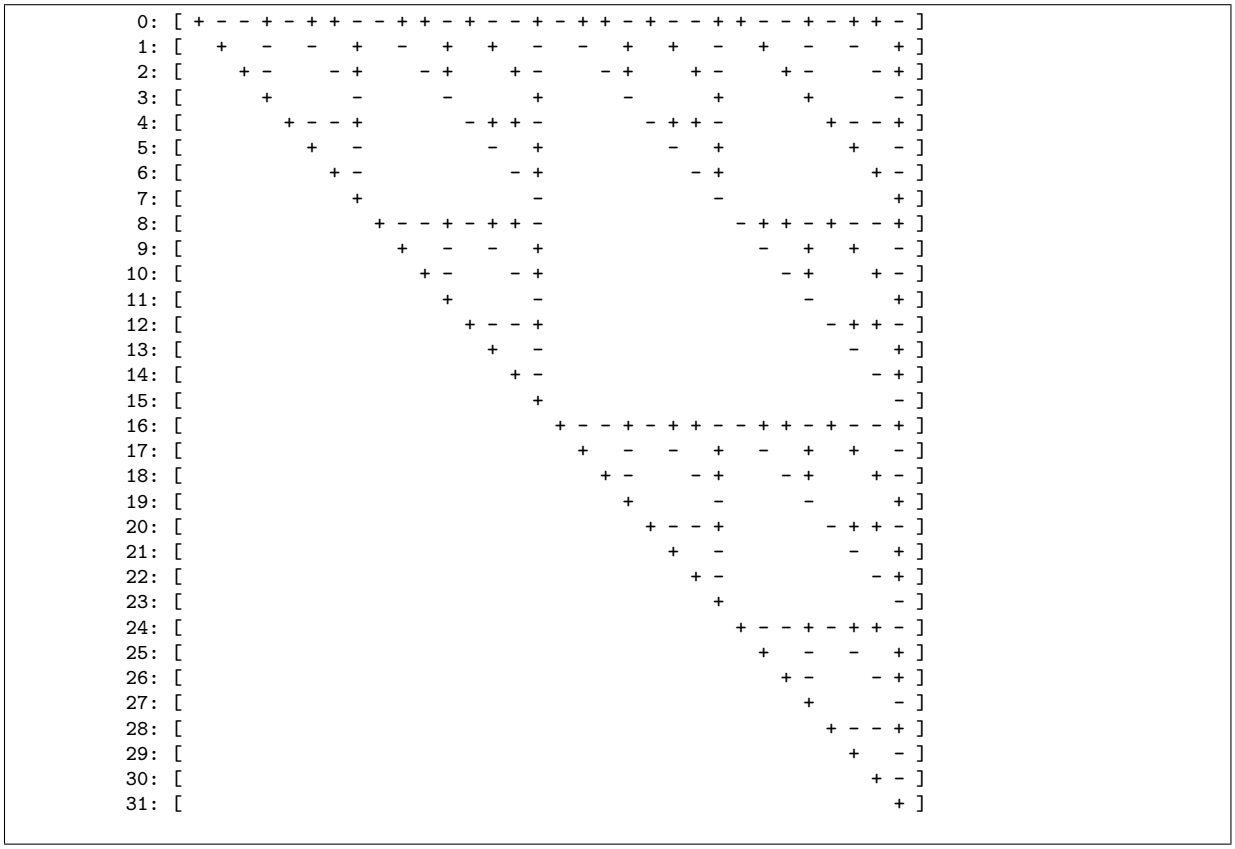

Figure 22.11-A: Basis functions for the Arithmetic transform (A^-, the one with the minus sign). The values are ± 1, blank entries denote 0.

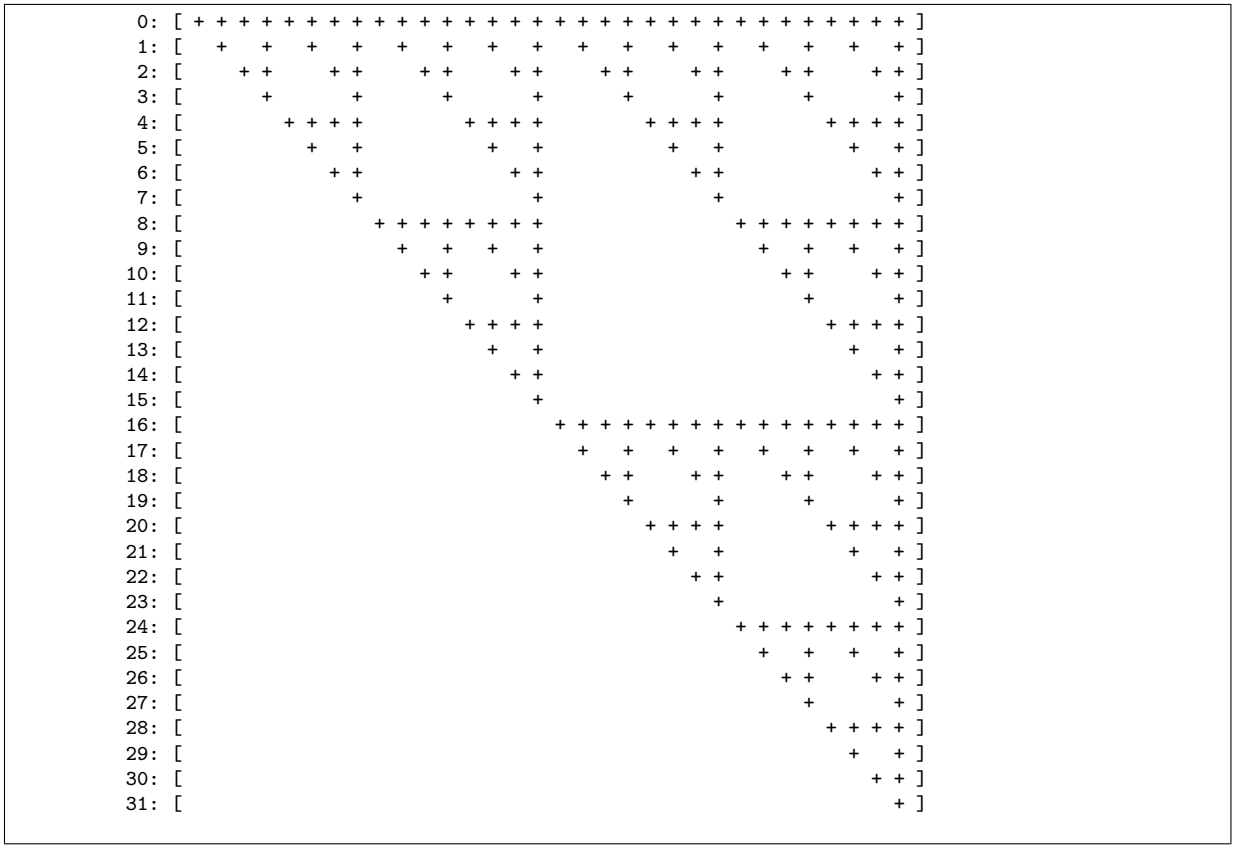

Figure 22.11-B: Basis functions for the inverse Arithmetic transform (A^+, the one without minus sign). The values are +1, blank entries denote 0.

On to the code [FXT: walsh/arithtransform.h]:

```
template <typename Type>
void arith_transform_plus(Type *f, ulong ldn)
// Arithmetic Transform (positive sign).
// Radix-2 decimation In Frequency (DIF) algorithm.
{
    const ulong n = (1UL<<ldn);
    for (ulong ldm=ldn; ldm>=1; --ldm)
    {
        const ulong m = (1UL<<ldm);
        const ulong mh = (m>>1);
        for (ulong r=0; r<n; r+=m)
        {
            ulong t1 = r;
            ulong t2 = r+mh;
            for (ulong j=0;  j<mh;  ++j, ++t1, ++t2)
            {
                Type u = f[t1];
                Type v = f[t2];
                f[t1] = u;
                f[t2] = u + v;
            }
        }
    }
}

```

and

```
template <typename Type>
void arith_transform_minus(Type *f, ulong ldn)
// Arithmetic Transform (negative sign).
```

```
4   // Radix-2 decimation In Frequency (DIF) algorithm.
5   // Inverse of arith_transform_plus().
6   {
7     [--snip--]
8                   f[t1] = u;
9                   f[t2] = v - u;
10    [--snip--]
11  }
```

The length-2 transforms can be written as

$$A_2^+ v = \begin{bmatrix} +1 & 0 \\ +1 & +1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ a+b \end{bmatrix} \tag{22.11-1}$$

$$A_2^- v = \begin{bmatrix} +1 & 0 \\ -1 & +1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ b-a \end{bmatrix} \tag{22.11-2}$$

That the transform with the minus is called the forward transform is tradition. Similar to the Fourier transform we avoid the forward- backward- naming scheme and put:

```
1   template <typename Type>
2   inline void arith_transform(Type *f, ulong ldn, int is)
3   {
4       if ( is>0 )  arith_transform_plus(f, ldn);
5       else         arith_transform_minus(f, ldn);
6   }
```

In Kronecker product notation (see section 22.3 on page 451) the arithmetic transform and its inverse can be written as

$$A_2^+ = \begin{bmatrix} +1 & 0 \\ +1 & +1 \end{bmatrix} \qquad A_n^+ = \bigotimes_{k=1}^{\log_2(n)} A_2^+ \tag{22.11-3a}$$

$$A_2^- = \begin{bmatrix} +1 & 0 \\ -1 & +1 \end{bmatrix} \qquad A_n^- = \bigotimes_{k=1}^{\log_2(n)} A_2^- \tag{22.11-3b}$$

The k-th element of the arithmetic A^+ transform is

$$A^+[a]_k = \sum_{i \subset k} a_i \tag{22.11-4a}$$

where $i \subset k$ means that the bits of i are a subset of the bits of k: $i \subset k \iff (i \wedge k) = i$. For the transform A^- we have

$$A^-[a]_k = (-1)^{p(k)} \sum_{i \subset k} (-1)^{p(i)} a_i \tag{22.11-4b}$$

where $p(x)$ is the parity of x.

22.11.1 Transposed arithmetic transform

We define a transposed arithmetic transforms B^+ and B^- via

$$B_2^+ = \begin{bmatrix} +1 & +1 \\ 0 & +1 \end{bmatrix} \qquad B_n^+ = \bigotimes_{k=1}^{\log_2(n)} B_2^+ \tag{22.11-5a}$$

$$B_2^- = \begin{bmatrix} +1 & -1 \\ 0 & +1 \end{bmatrix} \qquad B_n^- = \bigotimes_{k=1}^{\log_2(n)} B_2^- \tag{22.11-5b}$$

Then the transforms are [FXT: walsh/arithtransform.h]

```
1  template <typename Type>
2  void transposed_arith_transform_plus(Type *f, ulong ldn)
3  {
4    [--snip--]
5                  f[t1] = u + v;
6                  f[t2] = v;
7    [--snip--]
8  }
```

```
1  template <typename Type>
2  void transposed_arith_transform_minus(Type *f, ulong ldn)
3  {
4    [--snip--]
5                  f[t1] = u - v;
6                  f[t2] = v;
7    [--snip--]
8  }
```

22.11.2 Relation to Walsh transform

To establish the relation to the Walsh transform recall that its decomposition as a Kronecker product is

$$W_2 \;=\; \begin{bmatrix} +1 & +1 \\ +1 & -1 \end{bmatrix} \qquad W_n = \bigotimes_{k=1}^{\log_2(n)} W_2 \tag{22.11-6}$$

Now as $(W_2\,A_2^+)\,A_2^- = W_2$ the expression in parenthesis is the matrix that transforms the 2-point arithmetic transform (with the negative sign) to the Walsh transform. Similarly, as $(\frac{1}{2}\,A_2^+\,W_2)\,W_2 = A_2^+$, the matrix leading to the conversion from Walsh to arithmetic transform can be determined. One finds:

$$\left(W_2\,A_2^+\right)\,A_2^- \;=\; W_2 = \begin{bmatrix} +2 & +1 \\ 0 & -1 \end{bmatrix} A_2^- \tag{22.11-7a}$$

$$\left(W_2\,A_2^-\right)\,A_2^+ \;=\; W_2 = \begin{bmatrix} 0 & +1 \\ +2 & -1 \end{bmatrix} A_2^+ \tag{22.11-7b}$$

$$\left(\frac{1}{2}\,A_2^-\,W_2\right)\,W_2 \;=\; A_2^- = \frac{1}{2}\begin{bmatrix} +1 & +1 \\ 0 & -2 \end{bmatrix} W_2 \tag{22.11-7c}$$

$$\left(\frac{1}{2}\,A_2^+\,W_2\right)\,W_2 \;=\; A_2^+ = \frac{1}{2}\begin{bmatrix} +1 & +1 \\ +2 & 0 \end{bmatrix} W_2 \tag{22.11-7d}$$

The Kronecker product of the given matrices gives the converting transform. For example, using relation 22.11-7a, define

$$T_n \;:=\; \bigotimes_{k=1}^{\log_2(n)} \begin{bmatrix} +2 & +1 \\ 0 & -1 \end{bmatrix} \tag{22.11-8}$$

Then T_n converts a arithmetic (minus sign) transform to a Walsh transform: $W_n = T_n\,A_n^-$. For the relations between the arithmetic transform, the Reed-Muller transform and the Walsh transform, see [252].

22.12 Reed-Muller transform

How to make a Reed-Muller transform out of the arithmetic transform:
'*Replace* `+` *and* `-` *by XOR, done.*'

The *Reed-Muller transform* is can be obtained from the arithmetic transform by working modulo two. The transform is self-inverse, its basis functions are identical to those of the arithmetic transform A^+, shown in figure 22.11-B on page 474. The implementation is almost identical to [FXT: `walsh_wak_dif2()` in walsh/walshwak2.h]. The only changes are

```
Walsh: f[t1] = u + v;   =-->  Reed-Muller: f[t1] = u;
Walsh: f[t2] = u - v;   =-->  Reed-Muller: f[t2] = u ^ v;
```

There we go [FXT: walsh/reedmuller.h]:

```
template <typename Type>
void word_reed_muller_dif2(Type *f, ulong ldn)
// Reed-Muller Transform.
// Radix-2 decimation in frequency (DIF) algorithm.
// Self-inverse.
// Type must have the XOR operator.
{
    const ulong n = (1UL<<ldn);
    for (ulong ldm=ldn; ldm>=1; --ldm)
    {
        const ulong m = (1UL<<ldm);
        const ulong mh = (m>>1);
        for (ulong r=0; r<n; r+=m)
        {
            ulong t1 = r;
            ulong t2 = r+mh;
            for (ulong j=0;  j<mh;  ++j, ++t1, ++t2)
            {
                Type u = f[t1];
                Type v = f[t2];
                f[t1] = u;
                f[t2] = u ^ v;
            }
        }
    }
}
```

The decimation in time algorithm can be obtained from [FXT: `walsh_wak_dit2()` in walsh/walshwak2.h] by the very same changes. As given, the transforms work word-wise, if the bit-wise transform is wanted use

```
template <typename Type>
inline void bit_reed_muller(Type *f, ulong ldn)
{
    word_reed_muller_dif2(f, ldn);
    ulong n = 1UL << ldn;
    for (ulong k=0; k<n; ++k)  f[k] = yellow_code(f[k]);
}
```

The `yellow_code()` (see section 1.19 on page 49) can also be applied before the main loop. In fact, the yellow code *is* the Reed-Muller transform on a binary word.

The other 'color-transforms' of section 1.19 lead to variants of the Reed-Muller transform, the blue code gives another self-inverse transform, the red code and the green code give transforms R and E so that

$$R\,R\,R = \mathrm{id} \qquad R^{-1} = R\,R = E \tag{22.12-1}$$

$$E\,E\,E = \mathrm{id} \qquad E^{-1} = E\,E = R \tag{22.12-2}$$

$$R\,E = E\,R = \mathrm{id} \tag{22.12-3}$$

In fact, all relations given in the referenced section hold.

As can be seen from the 'atomic' matrices (relations 1.19-12c ... 1.19-12f on page 55) the four transforms corresponding to the 'color codes' are obtained by

```
Walsh:   f[t1] = u + v;    f[t2] = u - v;

   B:    f[t1] = u ^ v;    f[t2] = v;      (transposed Reed-Muller transform)
   Y:    f[t1] = u;        f[t2] = u ^ v;  (Reed-Muller transform)
   R:    f[t1] = v;        f[t2] = u ^ v;
   E:    f[t1] = u ^ v;    f[t2] = u;
```

The basis functions of the transforms are shown in figure 22.12-A.

The *transposed Reed-Muller transform* can be obtained by setting

```
Walsh: f[t1] = u + v;   =-->  transposed Reed-Muller: f[t1] = u ^ v;
Walsh: f[t2] = u - v;   =-->  transposed Reed-Muller: f[t2] = v;
```

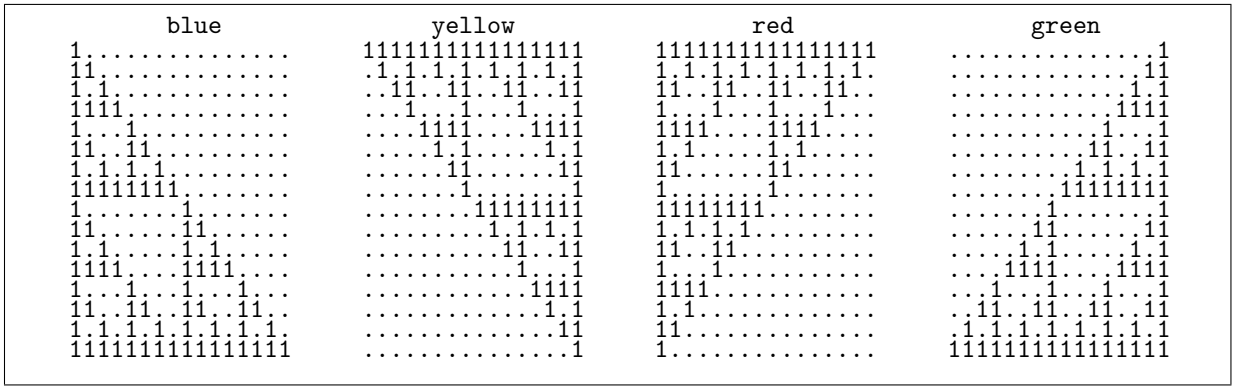

Figure 22.12-A: Basis functions of the length-16 blue, yellow, red, and green transforms.

The symbolic powering idea from section 1.19 on page 49 leads to transforms with bases (using eight element arrays and the yellow code):

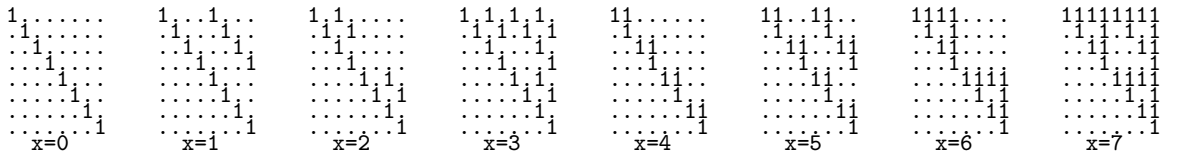

The program [FXT: bits/bitxtransforms-demo.cc] gives the matrices for 64-bit words.

A function that computes the k-th base function of the transform is [FXT: walsh/reedmuller.h]:

```
template <typename Type>
inline void reed_muller_basefunc(Type *f, ulong n, ulong k)
{
    for (ulong i=0; i<n; ++i)
    {
        f[i] = ( (i & k)==k ? +1 : 0 );  // is k a bit-subset of i ?
    }
}
```

Functions that are the word-wise equivalents of the Gray code are given in [FXT: aux1/wordgray.h]:

```
template <typename Type>
void word_gray(Type *f, ulong n)
{
    for (ulong k=0;  k<n-1;  ++k)  f[k] ^= f[k+1];
}
```

and

```
void inverse_word_gray(Type *f, ulong n)
{
    ulong x = 0,  k = n;
    while ( k-- )  { x ^= f[k];  f[k] = x; }
}
```

As one might suspect, these are related to the Reed-Muller transform. Writing Y ('yellow') for the Reed-Muller transform, g for the word-wise Gray code and S_k for the cyclic shift by k words (word zero is moved to position k) one has

$$Y\,S_{+1}\,Y = g \qquad (22.12\text{-}4)$$

$$Y\,S_{-1}\,Y = g^{-1} \qquad (22.12\text{-}5)$$

$$Y\,S_{k}\,Y = g^{k} \qquad (22.12\text{-}6)$$

These are exactly the relations 1.19-10a ... 1.19-10c on page 54 for the bit-wise transforms. For $k >= 0$ the operator S_k corresponds to the shift toward element zero (use [FXT: `rotate_sgn()` in perm/rotate.h]). The power of the word-wise Gray code is perfectly equivalent to the bit-wise version:

```
template <typename Type>
void word_gray_pow(Type *f, ulong n, ulong x)
{
    for (ulong s=1; s<n; s*=2)
    {
        if ( x & 1)
        {
            // word_gray ** s:
            for (ulong k=0, j=k+s;  j<n;  ++k,++j)  f[k] ^= f[j];
        }
        x >>= 1;
    }
}
```

Let e be the reversed Gray code operator, then we have for the transposed Reed-Muller transform B:

$$B\,S_{+1}\,B = e^{-1} \tag{22.12-7}$$
$$B\,S_{-1}\,B = e \tag{22.12-8}$$
$$B\,S_{k}\,B = e^{-k} \tag{22.12-9}$$

Further,

$$E\,S_{k}\,R = e^{k} \tag{22.12-10}$$
$$E\,e^{k}\,R = S_{k} \tag{22.12-11}$$

The transforms as Kronecker products (all operations are modulo two):

$$B_n = \bigotimes_{k=1}^{\log_2(n)} B_2 \quad \text{where} \quad B_2 = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \tag{22.12-12a}$$
$$Y_n = \bigotimes_{k=1}^{\log_2(n)} Y_2 \quad \text{where} \quad Y_2 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \tag{22.12-12b}$$
$$R_n = \bigotimes_{k=1}^{\log_2(n)} R_2 \quad \text{where} \quad R_2 = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} \tag{22.12-12c}$$
$$E_n = \bigotimes_{k=1}^{\log_2(n)} E_2 \quad \text{where} \quad E_2 = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \tag{22.12-12d}$$

22.13 The OR-convolution, and the AND-convolution

Let a and b be sequences of length a power of two. We define the *OR-convolution* h of a and b, as

$$h_\tau = \sum_{i \vee j = \tau} a_i\, b_j \tag{22.13-1}$$

where $\vee$ denotes bit-wise OR. The symbolic table for the OR-convolution is shown in figure 22.13-A (see figure 21.1-A on page 428 an explanation of the scheme). The OR-convolution can be computed as

$$h = A^{-}\left[\, A^{+}[a] \cdot A^{+}[b] \,\right] \tag{22.13-2}$$

where A^+ and A^- are the arithmetic transforms given in section 22.11 on page 473. An implementation is [FXT: walsh/or-convolution.h]:

```
template <typename Type>
inline void or_convolution(Type * restrict f, Type * restrict g, ulong ldn)
// Compute the OR-convolution h[] of f[] and g[]:
```

```
 +--  0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
 |
 0:   0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
 1:   1  1  3  3   5  5  7  7   9  9 11 11  13 13 15 15
 2:   2  3  2  3   6  7  6  7  10 11 10 11  14 15 14 15
 3:   3  3  3  3   7  7  7  7  11 11 11 11  15 15 15 15

 4:   4  5  6  7   4  5  6  7  12 13 14 15  12 13 14 15
 5:   5  5  7  7   5  5  7  7  13 13 15 15  13 13 15 15
 6:   6  7  6  7   6  7  6  7  14 15 14 15  14 15 14 15
 7:   7  7  7  7   7  7  7  7  15 15 15 15  15 15 15 15

 8:   8  9 10 11  12 13 14 15   8  9 10 11  12 13 14 15
 9:   9  9 11 11  13 13 15 15   9  9 11 11  13 13 15 15
10:  10 11 10 11  14 15 14 15  10 11 10 11  14 15 14 15
11:  11 11 11 11  15 15 15 15  11 11 11 11  15 15 15 15

12:  12 13 14 15  12 13 14 15  12 13 14 15  12 13 14 15
13:  13 13 15 15  13 13 15 15  13 13 15 15  13 13 15 15
14:  14 15 14 15  14 15 14 15  14 15 14 15  14 15 14 15
15:  15 15 15 15  15 15 15 15  15 15 15 15  15 15 15 15
```

Figure 22.13-A: Semi-symbolic scheme for the OR-convolution of two length-16 sequences.

```
 +--  0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
 |
 0:   0  0  0  0   0  0  0  0   0  0  0  0   0  0  0  0
 1:   0  1  0  1   0  1  0  1   0  1  0  1   0  1  0  1
 2:   0  0  2  2   0  0  2  2   0  0  2  2   0  0  2  2
 3:   0  1  2  3   0  1  2  3   0  1  2  3   0  1  2  3

 4:   0  0  0  0   4  4  4  4   0  0  0  0   4  4  4  4
 5:   0  1  0  1   4  5  4  5   0  1  0  1   4  5  4  5
 6:   0  0  2  2   4  4  6  6   0  0  2  2   4  4  6  6
 7:   0  1  2  3   4  5  6  7   0  1  2  3   4  5  6  7

 8:   0  0  0  0   0  0  0  0   8  8  8  8   8  8  8  8
 9:   0  1  0  1   0  1  0  1   8  9  8  9   8  9  8  9
10:   0  0  2  2   0  0  2  2   8  8 10 10   8  8 10 10
11:   0  1  2  3   0  1  2  3   8  9 10 11   8  9 10 11

12:   0  0  0  0   4  4  4  4   8  8  8  8  12 12 12 12
13:   0  1  0  1   4  5  4  5   8  9  8  9  12 13 12 13
14:   0  0  2  2   4  4  6  6   8  8 10 10  12 12 14 14
15:   0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
```

Figure 22.13-B: Semi-symbolic scheme for the AND-convolution of two length-16 sequences.

```
4   //  h[k] = sum(i | j == k,  f[i]*g[j])
5   // f[] and g[] must not overlap.
6   // Result written to g[].
7   {
8       arith_transform_plus(f, ldn);
9       arith_transform_plus(g, ldn);
10      const ulong n = (1UL<<ldn);
11      for (ulong k=0; k<n; ++k)  g[k] *= f[k];
12      arith_transform_minus(g, ldn);
13  }
```

We also have

$$h \;=\; Y^{-1}\,[\,Y[a]\cdot Y[b]\,] \;=\; Y\,[\,Y[a]\cdot Y[b]\,] \tag{22.13-3}$$

where Y is the Reed-Muller transform given in section 22.12 on page 476.

Define the *AND-convolution* h of two sequences a and b as

$$h_\tau \;=\; \sum_{i\wedge j=\tau} a_i\, b_j \tag{22.13-4}$$

where $\wedge$ denotes the bit-wise AND. The symbolic scheme is shown in figure 22.13-B. The AND-convolution can be computed as

$$h \;=\; B^-\,\big[\,B^+[a]\cdot B^+[b]\,\big] \tag{22.13-5}$$

where B^+ and B^- are the transposed arithmetic transforms. The implementation of the AND-convolution is [FXT: walsh/and-convolution.h]:

```
template <typename Type>
inline void and_convolution(Type * restrict f, Type * restrict g, ulong ldn)
// Compute the AND-convolution h[] of f[] and g[]:
//  h[k] = sum(i & j == k,  f[i]*g[j])
// f[] and g[] must not overlap.
// Result written to g[].
{
    transposed_arith_transform_plus(f, ldn);
    transposed_arith_transform_plus(g, ldn);
    const ulong n = (1UL<<ldn);
    for (ulong k=0; k<n; ++k)  g[k] *= f[k];
    transposed_arith_transform_minus(g, ldn);
}
```

The same is true modulo two:

$$h \;=\; B^{-1}\,[\,B[a]\cdot B[b]\,] \;=\; B\,[\,B[a]\cdot B[b]\,] \tag{22.13-6}$$

here B is the transposed Reed-Muller transform.

Chapter 23

The Haar transform

Haar transforms are invertible transforms that do not involve trigonometric factors. We present several variants of the transform whose computation involves just proportional n operations. Haar transforms can be used as building blocks of the Walsh transform. We give two non-standard splitting schemes for Haar transforms, based on the Fibonacci and Mersenne numbers.

23.1 The 'standard' Haar transform

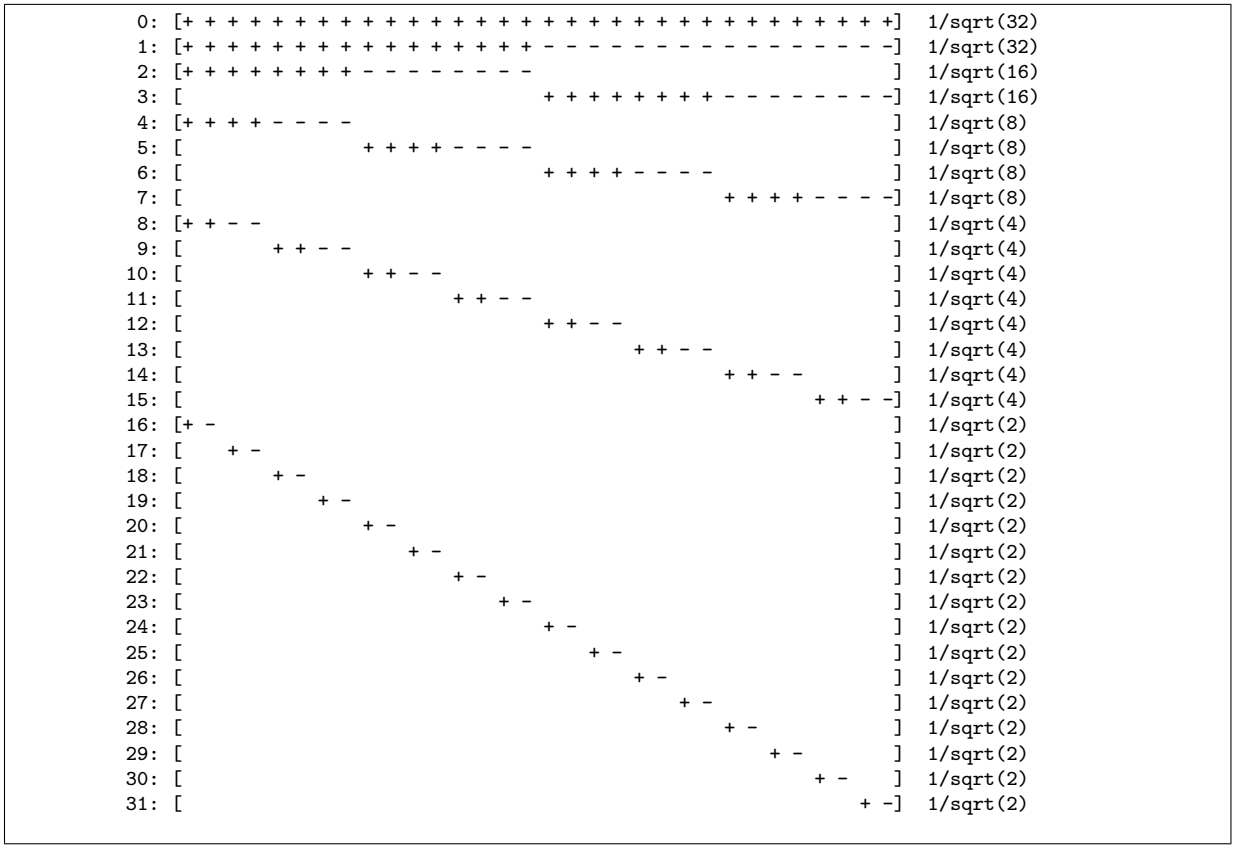

Figure 23.1-A: Basis functions for the Haar transform. Only the signs of the nonzero entries are shown. The absolute value of the nonzero entries in each row is given at the right. The norm of each row is one.

The *Haar transform* of a length-n sequence f consists of $\log_2(n)$ steps were the sums and differences of adjacent pairs of elements f_{2j}, f_{2j+1} are computed. The sums are then written to the lower half of the array f, the differences to the higher half. Ignoring the order (and normalization), each step corresponds to a matrix multiplication:

$$\begin{bmatrix} +1 & +1 & & & & & & & \\ +1 & -1 & & & & & & & \\ & & +1 & +1 & & & & & \\ & & +1 & -1 & & & & & \\ & & & & +1 & +1 & & & \\ & & & & +1 & -1 & & & \\ & & & & & & +1 & +1 & \\ & & & & & & +1 & -1 & \\ & & & & & & & & \ddots \end{bmatrix} \cdot \begin{bmatrix} f_0 \\ f_1 \\ f_2 \\ f_3 \\ f_4 \\ f_5 \\ f_6 \\ f_7 \\ \vdots \end{bmatrix} \tag{23.1-1}$$

The step is applied to the full array, then to the lower half, the lower quarter, ..., the lower four elements, the lowest pair. (The array length n must be a power of two.) The computational cost of the transform is proportional $n + n/2 + n/4 + \ldots + 4 + 2$ which is $\sim O(n)$. The basis functions for the Haar transform have finite support, they are shown in figure 23.1-A.

The following implementation involves $2n$ multiplications $\sqrt{2}$ which make the transform orthogonal, corresponding to a scalar factor of $\sqrt{2}$ in the relation 23.1-1:

```
template <typename Type>
void haar(Type *f, ulong ldn)
{
    ulong n = (1UL<<ldn);
    const Type s2 = sqrt(0.5);   // normalization factor
    Type *g = new Type[n];       // scratch space
    for (ulong m=n; m>1; m>>=1)  // n, n/2, n/4, n/8, ..., 4, 2
    {
        ulong mh = (m>>1);
        for (ulong j=0, k=0;  j<m;  j+=2, k++)  // sums and differences of adjacent pairs
        {
            Type x = f[j];
            Type y = f[j+1];
            g[k]    = (x + y) * s2;   // sums to lower half
            g[mh+k] = (x - y) * s2;   // differences to higher half
        }
        copy(g, f, m);
    }
    delete [] g;
}
```

We can reduce the number of multiplications to n by delaying the multiplies with the sums [FXT: `haar()` in haar/haar.h]:

```
template <typename Type>
void haar(Type *f, ulong ldn, Type *ws=0)
{
    ulong n = (1UL<<ldn);
    Type s2 = sqrt(0.5);
    Type v = 1.0;
    Type *g = ws;
    if ( !ws )  g = new Type[n];
    for (ulong m=n; m>1; m>>=1)
    {
        v *= s2;
        ulong mh = (m>>1);
        for (ulong j=0, k=0;  j<m;  j+=2, k++)
        {
            Type x = f[j];
            Type y = f[j+1];
            g[k]    =  x + y;
            g[mh+k] = (x - y) * v;
```

```
19          }
20          copy(g, f, m);
21      }
22      f[0] *= v; // v == 1.0/sqrt(n);
23      if ( !ws )  delete [] g;
24  }
```

The temporary workspace can be supplied by the caller.

The inverse Haar transform is obtained by applying the inverse steps in reversed order [FXT: `inverse_haar()` in haar/haar.h]:

```
1   template <typename Type>
2   void inverse_haar(Type *f, ulong ldn, Type *ws=0)
3   {
4       ulong n = (1UL<<ldn);
5       Type s2 = sqrt(2.0);
6       Type v = 1.0/sqrt(n);
7       Type *g = ws;
8       if ( !ws )  g = new Type[n];
9       f[0] *= v;
10      for (ulong m=2; m<=n; m<<=1)
11      {
12          ulong mh = (m>>1);
13          for (ulong j=0, k=0;  j<m;  j+=2, k++)
14          {
15              Type x = f[k];
16              Type y = f[mh+k] * v;
17              g[j]    =  x + y;
18              g[j+1]  =  x - y;
19          }
20          copy(g, f, m);
21          v *= s2;
22      }
23      if ( !ws )  delete [] g;
24  }
```

A generalization of the steps used in the Haar transform leads to the wavelet transforms treated in chapter 26 on page 533.

23.2 In-place Haar transform

The 'standard' Haar transform routines are not in-place, they use a temporary storage. A rather simple reordering of the basis functions, however, allows for to an in-place algorithm [FXT: `haar_inplace()` in haar/haar.h]:

```
1   template <typename Type>
2   void haar_inplace(Type *f, ulong ldn)
3   {
4       ulong n = 1UL<<ldn;
5       Type s2 = sqrt(0.5);
6       Type v = 1.0;
7       for (ulong js=2; js<=n; js<<=1)
8       {
9           v *= s2;
10          for (ulong j=0, t=js>>1;  j<n;  j+=js, t+=js)
11          {
12              Type x = f[j];
13              Type y = f[t];
14              f[j]  =  x + y;
15              f[t]  = (x - y) * v;
16          }
17      }
18      f[0] *= v; // v==1.0/sqrt(n);
19  }
```

The inverse is [FXT: `inverse_haar_inplace()` in haar/haar.h]:

```
1   template <typename Type>
```

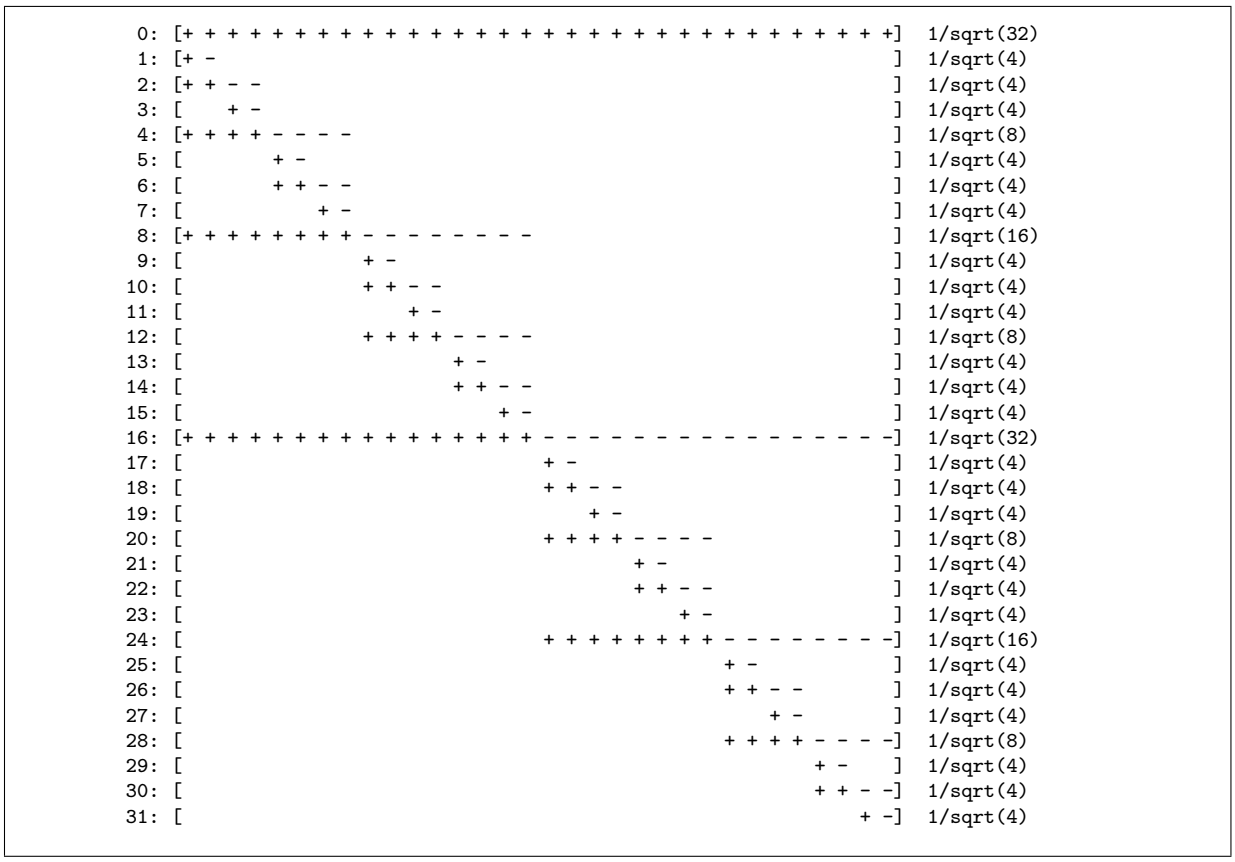

Figure 23.2-A: Haar basis functions, in-place order. Only the signs of the nonzero entries are shown. The absolute value of the nonzero entries in each row is given at the right. The norm of each row is one.

```
2   void inverse_haar_inplace(Type *f, ulong ldn)
3   {
4       ulong n = 1UL<<ldn;
5       Type s2 = sqrt(2.0);
6       Type v = 1.0/sqrt(n);
7       f[0] *= v;
8       for (ulong js=n; js>=2; js>>=1)
9       {
10          for (ulong j=0, t=js>>1;  j<n;  j+=js, t+=js)
11          {
12              Type x = f[j];
13              Type y = f[t] * v;
14              f[j]  =  x + y;
15              f[t]  =  x - y;
16          }
17          v *= s2;
18      }
19  }
```

The in-place Haar transform H_i is related to the 'usual' Haar transform H by a permutation P_H via the relations

$$H = P_H \cdot H_i \tag{23.2-1}$$

$$H^{-1} = H_i^{-1} \cdot P_H^{-1} \tag{23.2-2}$$

The permutation P_H can be programmed as

```
1   template <typename Type>
2   void haar_permute(Type *f, ulong n)
3   {
```

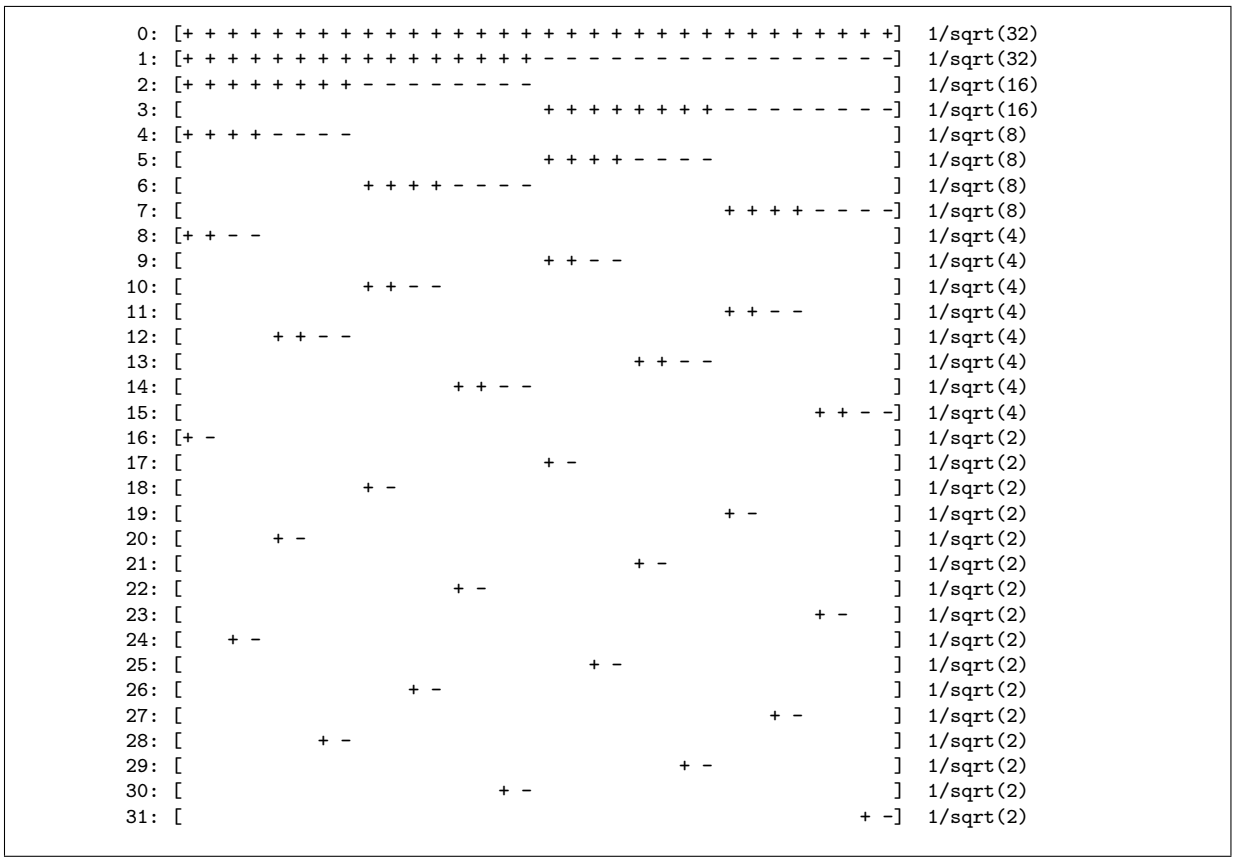

Figure 23.2-B: Basis functions of the in-place order Haar transform followed by a revbin permutation. The ordering is such that basis functions that are identical up to a shift appear consecutively.

```
4       revbin_permute(f, n);
5       for (ulong m=4; m<=n/2; m*=2)  revbin_permute(f+m, m);
6   }
```

The revbin permutations in the loop do not overlap, so the inverse Haar permutation is obtained by simply swapping the loop with the full-length revbin permutation [FXT: perm/haarpermute.h]:

```
1   template <typename Type>
2   void inverse_haar_permute(Type *f, ulong n)
3   {
4       for (ulong m=4; m<=n/2; m*=2)  revbin_permute(f+m, m);
5       revbin_permute(f, n);
6   }
```

Then, as given above, `haar` is equivalent to

```
haar_inplace();  haar_permute();
```

and `inverse_haar` is equivalent to

```
inverse_haar_permute();  inverse_haar_inplace();
```

23.3 Non-normalized Haar transforms

Versions of the Haar transform without normalization are given in [FXT: haar/haarnn.h]. The basis functions are the same as for the normalized versions, only the absolute value of the nonzero entries are different.

```
1   template <typename Type>
2   void haar_nn(Type *f, ulong ldn, Type *ws=0)
```

```
{
    ulong n = (1UL<<ldn);
    Type *g = ws;
    if ( !ws )  g = new Type[n];
    for (ulong m=n; m>1; m>>=1)
    {
        ulong mh = (m>>1);
        for (ulong j=0, k=0;  j<m;  j+=2, k++)
        {
            Type x = f[j];
            Type y = f[j+1];
            g[k]    =  x + y;
            g[mh+k] =  x - y;
        }
        copy(g, f, m);
    }
    if ( !ws )  delete [] g;
}
```

The inverse is

```
template <typename Type>
void inverse_haar_nn(Type *f, ulong ldn, Type *ws=0)
{
    ulong n = (1UL<<ldn);
    Type s2 = 2.0;
    Type v = 1.0/n;
    Type *g = ws;
    if ( !ws )  g = new Type[n];
    f[0] *= v;
    for (ulong m=2; m<=n; m<<=1)
    {
        ulong mh = (m>>1);

        for (ulong j=0, k=0;  j<m;  j+=2, k++)
        {
            Type x = f[k];
            Type y = f[mh+k] * v;
            g[j]    =  x + y;
            g[j+1]  =  x - y;
        }
        copy(g, f, m);
        v *= s2;
    }
    if ( !ws )  delete [] g;
}
```

An unnormalized transform that works in-place is

```
template <typename Type>
void haar_inplace_nn(Type *f, ulong ldn)
{
    ulong n = 1UL<<ldn;
    for (ulong js=2; js<=n; js<<=1)
    {
        for (ulong j=0, t=js>>1;  j<n;  j+=js, t+=js)
        {
            Type x = f[j];
            Type y = f[t];
            f[j]  =  x + y;
            f[t]  =  x - y;
        }
    }
}
```

The inverse routine is

```
template <typename Type>
void inverse_haar_inplace_nn(Type *f, ulong ldn)
{
    ulong n = 1UL<<ldn;
    Type s2 = 2.0;
    Type v = 1.0/n;
    f[0] *= v;
    for (ulong js=n; js>=2; js>>=1)
```

```
9       {
10          for (ulong j=0, t=js>>1;  j<n;  j+=js, t+=js)
11          {
12              Type x = f[j];
13              Type y = f[t] * v;
14              f[j]  =  x + y;
15              f[t]  =  x - y;
16          }
17          v *= s2;
18      }
19  }
```

The sequence of statements `{ haar_inplace_nn(); haar_permute(); }` is equivalent to `{ haar_nn(); }`. The sequence `{ inverse_haar_permute(); inverse_haar_inplace(); }` is equivalent to `{ inverse_haar(); }`.

23.4 Transposed Haar transforms

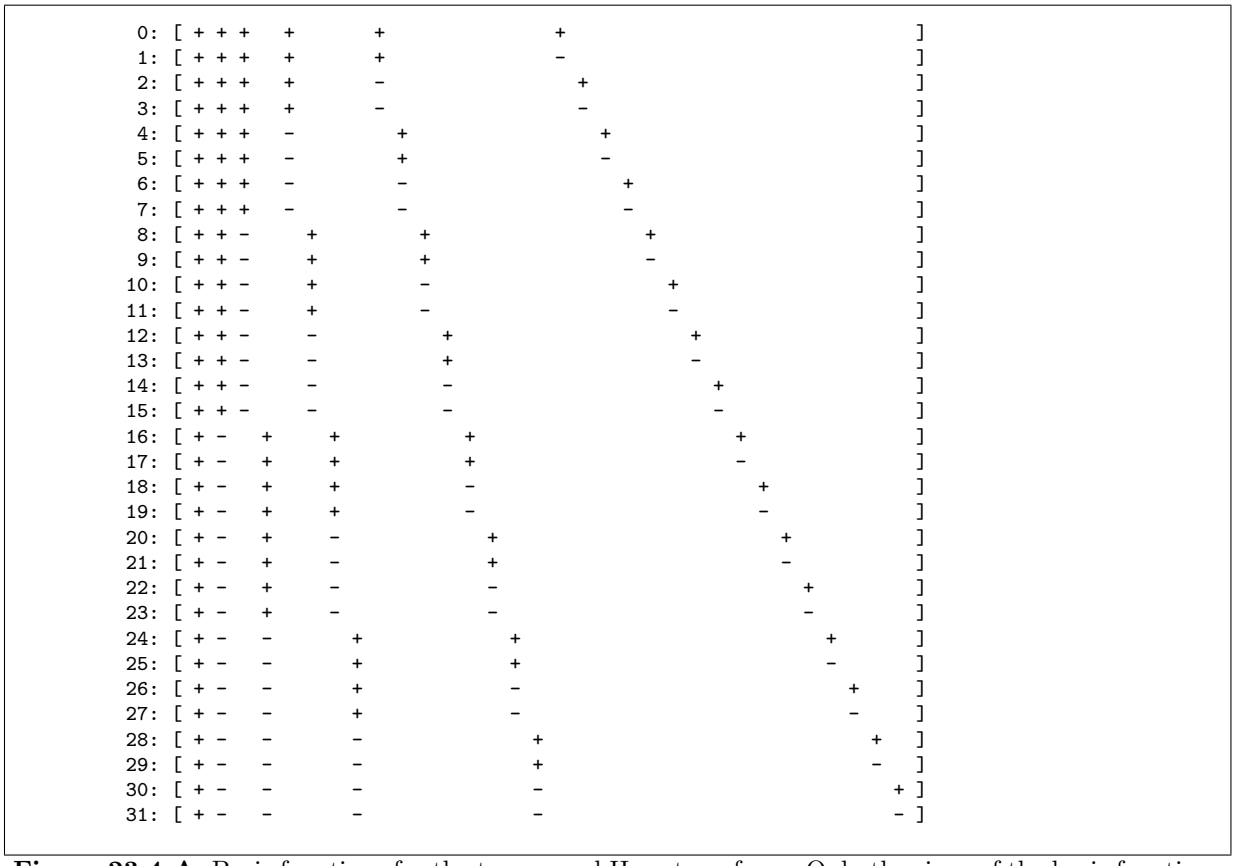

Figure 23.4-A: Basis functions for the transposed Haar transform. Only the signs of the basis functions are shown. At the blank entries the functions are zero.

Figure 23.4-A shows the bases functions of the transposed Haar transform. The shown routine are given in [FXT: haar/transposedhaarnn.h]. The following routine does an unnormalized Haar transform. The result is, up to normalization, the same as with `inverse_haar()`. The implementation uses a scratch array:

```
1   template <typename Type>
2   void transposed_haar_nn(Type *f, ulong ldn, Type *ws=0)
3   {
4       ulong n = (1UL<<ldn);
5       Type *g = ws;
```

```
6       if ( !ws )  g = new Type[n];
7       for (ulong m=2; m<=n; m<<=1)
8       {
9           ulong mh = (m>>1);
10          for (ulong j=0, k=0;  j<m;  j+=2, k++)
11          {
12              Type x = f[k];
13              Type y = f[mh+k];
14              g[j]    =  x + y;
15              g[j+1]  =  x - y;
16          }
17          copy(g, f, m);
18      }
19      if ( !ws )  delete [] g;
20  }
```

The inverse transform is

```
1   template <typename Type>
2   void inverse_transposed_haar_nn(Type *f, ulong ldn, Type *ws=0)
3   {
4       ulong n = (1UL<<ldn);
5       Type *g = ws;
6       if ( !ws )  g = new Type[n];
7       for (ulong m=n; m>1; m>>=1)
8       {
9           ulong mh = (m>>1);
10          for (ulong j=0, k=0;  j<m;  j+=2, k++)
11          {
12              Type x = f[j]   * 0.5;
13              Type y = f[j+1] * 0.5;
14              g[k]    =  x + y;
15              g[mh+k] =  x - y;
16          }
17          copy(g, f, m);
18      }
19      if ( !ws )  delete [] g;
20  }
```

The next routine is equivalent to the sequence of statements `{ inverse_haar_permute(); transposed_haar_inplace_nn(); }`. Its advantage is that no scratch array is needed:

```
1   template <typename Type>
2   void transposed_haar_inplace_nn(Type *f, ulong ldn)
3   {
4       ulong n = 1UL<<ldn;
5       for (ulong js=n; js>=2; js>>=1)
6       {
7           for (ulong j=0, t=js>>1;  j<n;  j+=js, t+=js)
8           {
9               Type x = f[j];
10              Type y = f[t];
11              f[j]  =  x + y;
12              f[t]  =  x - y;
13          }
14      }
15  }
```

The inverse transform is

```
1   template <typename Type>
2   void inverse_transposed_haar_inplace_nn(Type *f, ulong ldn)
3   {
4       ulong n = 1UL<<ldn;
5       for (ulong js=2; js<=n; js<<=1)
6       {
7           for (ulong j=0, t=js>>1;  j<n;  j+=js, t+=js)
8           {
9               Type x = f[j] * 0.5;
10              Type y = f[t] * 0.5;
11              f[j]  =  x + y;
12              f[t]  =  x - y;
13          }
14      }
15  }
```

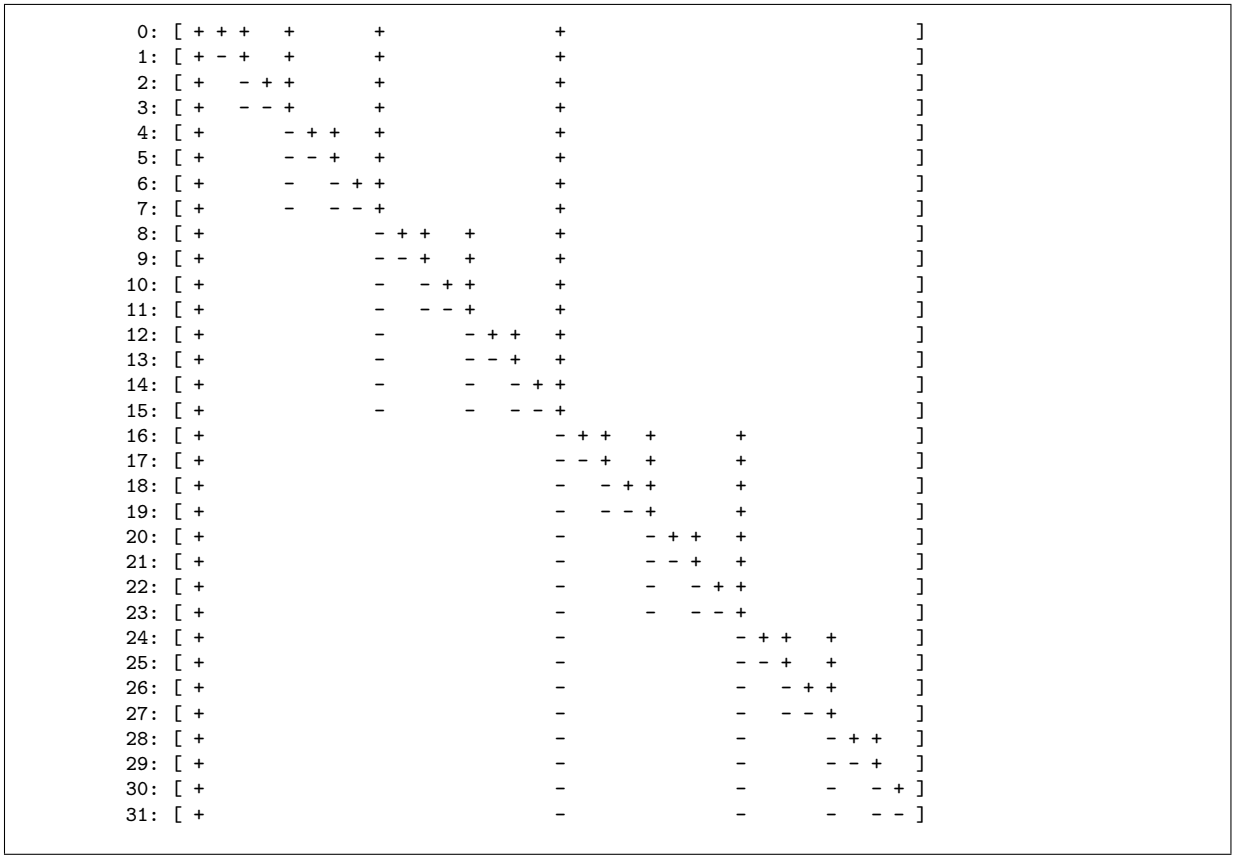

Figure 23.4-B: Basis functions for the transposed in-place Haar transform. Only the signs of the basis functions are shown. At the blank entries the functions are zero.

23.5 The reversed Haar transform

Let H_{ni} denote the non-normalized in-place Haar transform (`haar_inplace_nn`), Let H_{tni} denote the transposed non-normalized in-place Haar transform (`transposed_haar_inplace_nn`), R the revbin permutation, $\overline{H}$ the reversed Haar transform and $\overline{H}_t$ the transposed reversed Haar transform. Then

$$\overline{H} = R\,H_{ni}\,R \tag{23.5-1a}$$
$$\overline{H}_t = R\,H_{tni}\,R \tag{23.5-1b}$$
$$\overline{H}^{-1} = R\,H_{ni}^{-1}\,R \tag{23.5-1c}$$
$$\overline{H}_t^{-1} = R\,H_{tni}^{-1}\,R \tag{23.5-1d}$$

Code for the reversed Haar transform [FXT: `haar_rev_nn()` in haar/haarrevnn.h]:

```
1   template <typename Type>
2   void haar_rev_nn(Type *f, ulong ldn)
3   {
4   //    const ulong n = (1UL<<ldn);
5       for (ulong ldm=ldn; ldm>=1; --ldm)
6       {
7           const ulong m = (1UL<<ldm);
8           const ulong mh = (m>>1);
9           ulong r = 0;
10  //        for (ulong r=0; r<n; r+=m)  // almost walsh_wak_dif2()
11          {
12              ulong t1 = r;
13              ulong t2 = r + mh;
14              for (ulong j=0;  j<mh;  ++j, ++t1, ++t2)
```

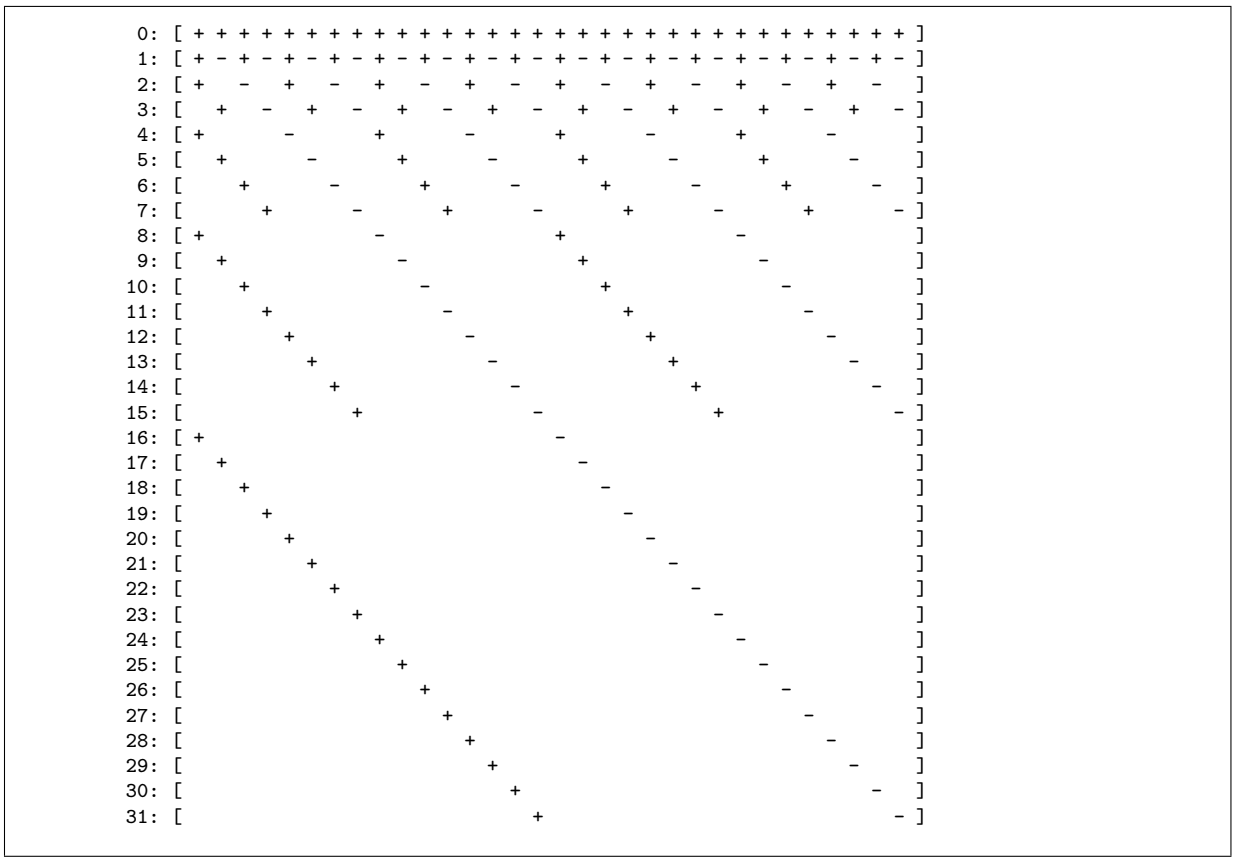

Figure 23.5-A: Basis functions for the reversed Haar transform. Only the signs of the nonzero entries are shown.

```
15              {
16                  Type u = f[t1];
17                  Type v = f[t2];
18                  f[t1] = u + v;
19                  f[t2] = u - v;
20              }
21          }
22      }
23  }
```

Note that this is almost the radix-2 DIF implementation for the Walsh transform The only change is that the line `for (ulong r=0; r<n; r+=m)` was replaced by `ulong r = 0`. The transform can also be computed via the following sequence of statements: `{ revbin_permute(); haar_inplace_nn(); revbin_permute(); }`.

The inverse transform is obtained by the equivalent modification with the DIT implementation for the Walsh transform and normalization:

```
1   template <typename Type>
2   void inverse_haar_rev_nn(Type *f, ulong ldn)
3   {
4       for (ulong ldm=1; ldm<=ldn; ++ldm)
5       {
6           const ulong m = (1UL<<ldm);
7           const ulong mh = (m>>1);
8           ulong r = 0;
9   //        for (ulong r=0; r<n; r+=m)  // almost walsh_wak_dit2()
10          {
11              ulong t1 = r;
12              ulong t2 = r + mh;
13              for (ulong j=0;  j<mh;  ++j, ++t1, ++t2)
14              {
```

```
15                  Type u = f[t1] * 0.5;
16                  Type v = f[t2] * 0.5;
17                  f[t1] = u + v;
18                  f[t2] = u - v;
19              }
20          }
21      }
22  }
```

The reversed transposed Haar transform is, up to normalization, the inverse of `haar_rev_nn()`. It is given in [FXT: haar/transposedhaarrevnn.h]:

```
1   template <typename Type>
2   void transposed_haar_rev_nn(Type *f, ulong ldn)
3   {
4       for (ulong ldm=1; ldm<=ldn; ++ldm)
5       {
6           const ulong m = (1UL<<ldm);
7           const ulong mh = (m>>1);
8           ulong r = 0;
9   //        for (ulong r=0; r<n; r+=m)  // almost walsh_wak_dit2()
10          {
11              ulong t1 = r;
12              ulong t2 = r + mh;
13              for (ulong j=0;  j<mh;  ++j, ++t1, ++t2)
14              {
15                  Type u = f[t1];
16                  Type v = f[t2];
17                  f[t1] = u + v;
18                  f[t2] = u - v;
19              }
20          }
21      }
22  }
```

The same result would be obtained by the following sequence of statements: `{ revbin_permute(); transposed_haar_inplace_nn(); revbin_permute(); }`. The inverse transform is

```
1   template <typename Type>
2   void inverse_transposed_haar_rev_nn(Type *f, ulong ldn)
3   {
4   //    const ulong n = (1UL<<ldn);
5       for (ulong ldm=ldn; ldm>=1; --ldm)
6       {
7           const ulong m = (1UL<<ldm);
8           const ulong mh = (m>>1);
9           ulong r = 0;
10  //        for (ulong r=0; r<n; r+=m)  // almost walsh_wak_dif2()
11          {
12              ulong t1 = r;
13              ulong t2 = r + mh;
14              for (ulong j=0;  j<mh;  ++j, ++t1, ++t2)
15              {
16                  Type u = f[t1] * 0.5;
17                  Type v = f[t2] * 0.5;
18                  f[t1] = u + v;
19                  f[t2] = u - v;
20              }
21          }
22      }
23  }
```

23.6 Relations between Walsh and Haar transforms

23.6.1 Walsh transforms from Haar transforms

A length-n Walsh transform can be obtained from one length-n Haar transform, one transform of length-$\frac{n}{2}$, two transforms of length-$\frac{n}{4}$, four transforms of length-$\frac{n}{8}$, ... and $\frac{n}{4}$ transforms of length-2. Using the reversed Haar transform the implementation is most straightforward: A Walsh transform (W_k, the one with the Walsh Kronecker base) can be implemented as

```
Haar transforms:
    H(16)                 H(8)          H(4)      H(2)
  AAAAAAAAaaaaaaaa        BBBBbbbb      CCcc      Dd
  AAAAaaaa                BBbb          Cc
  AAaa                    Bb
  Aa

Walsh(16) =^= 1*H(16) + 1*H(8) + 2*H(4) + 4*H(2)
  AAAAAAAAaaaaaaaa
  AAAAaaaaBBBBbbbb
  AAaaCCccBBbbCCcc
  AaDdCcDdBbDdCcDd
```

Figure 23.6-A: Symbolic description of how to build a Walsh transform from Haar transforms.

```
Transposed Haar transforms:
    H(16)                 H(8)          H(4)      H(2)
  Aa
  AAaa                    Bb
  AAAAaaaa                BBbb          Cc
  AAAAAAAAaaaaaaaa        BBBBbbbb      CCcc      Dd

Walsh(16) =^= 1*H(16) + 1*H(8) + 2*H(4) + 4*H(2)
  AaDdCcDdBbDdCcDd
  AAaaCCccBBbbCCcc
  AAAAaaaaBBBBbbbb
  AAAAAAAAaaaaaaaa
```

Figure 23.6-B: Symbolic description of how to build a Walsh transform from Haar transforms, transposed version.

```
1       // algorithm WH1:
2       ulong n = 1UL<<ldn;
3       haar_rev_nn(f, ldn);
4       for (ulong ldk=ldn-1; ldk>0; --ldk)
5       {
6           ulong k = 1UL << ldk;
7           for (ulong j=k; j<n; j+=2*k)  haar_rev_nn(f+j, ldk);
8       }
```

The idea, as a symbolic scheme, is shown in figure 23.6-A. The scheme obtained by reversing the order of the lines is shown in figure 23.6-B. It corresponds to the computation of W_k using the transposed version of the Haar transform:

```
1       // algorithm WH1T:
2       ulong n = 1UL<<ldn;
3       for (ulong ldk=1; ldk<ldn; ++ldk)
4       {
5           ulong k = 1UL << ldk;
6           for (ulong j=k; j<n; j+=2*k)  transposed_haar_rev_nn(f+j, ldk);
7       }
8       transposed_haar_rev_nn(f, ldn);
```

Two more methods are obtained by reversing the individual lines of the schemes seen so far, see figure 23.6-C. These correspond to the computation of the inverse Walsh transform ($W_k^{-1} = \frac{1}{n} W_k$) either as

```
1       // algorithm WH2T:
2       ulong n = 1UL<<ldn;
3       inverse_transposed_haar_rev_nn(f, ldn);
4       for (ulong ldk=ldn-1; ldk>0; --ldk)
```

Figure 23.6-C: Symbolic scheme of the four versions of the computation of the Walsh transform via Haar transforms.

```
5       {
6           ulong k = 1UL << ldk;
7           for (ulong j=k; j<n; j+=2*k)  inverse_transposed_haar_rev_nn(f+j, ldk);
8       }
```

or as

```
1       // algorithm WH2:
2       ulong n = 1UL<<ldn;
3       for (ulong ldk=1; ldk<ldn; ++ldk)
4       {
5           ulong k = 1UL << ldk;
6           for (ulong j=k; j<n; j+=2*k)  inverse_haar_rev_nn(f+j, ldk);
7       }
8       inverse_haar_rev_nn(f, ldn);
```

23.6.2 Haar transforms from Walsh transforms

The ($\sim n \log(n)$) schemes given here are not a efficient method to compute the Haar transform (which is $\sim n$). Instead, they can be used to identify the type of Haar transform that is the building block of a given Walsh transform.

The non-normalized transposed reversed Haar transform can (up to normalization) be obtained via

```
1       // algorithm HW1:  transposed_haar_rev_nn(f, ldn); =^=
2       for (ulong ldk=1; ldk<ldn; ++ldk)
3       {
4           ulong k = 1UL << ldk;
5           walsh_wak(f+k, ldk);
6       }
7       walsh_wak(f, ldn);
```

and its inverse as

```
1       // algorithm HW1I:  inverse_transposed_haar_rev_nn(f, ldn); =^=
2       walsh_wak(f, ldn);
3       for (ulong ldk=1; ldk<ldn; ++ldk)
4       {
5           ulong k = 1UL << ldk;
6           walsh_wak(f+k, ldk);
7       }
```

The non-normalized transposed Haar transform can (again, up to normalization) be obtained via

```
1       // algorithm HW2:  transposed_haar_nn(f, ldn); =^=
2       for (ulong ldk=1; ldk<ldn; ++ldk)
```

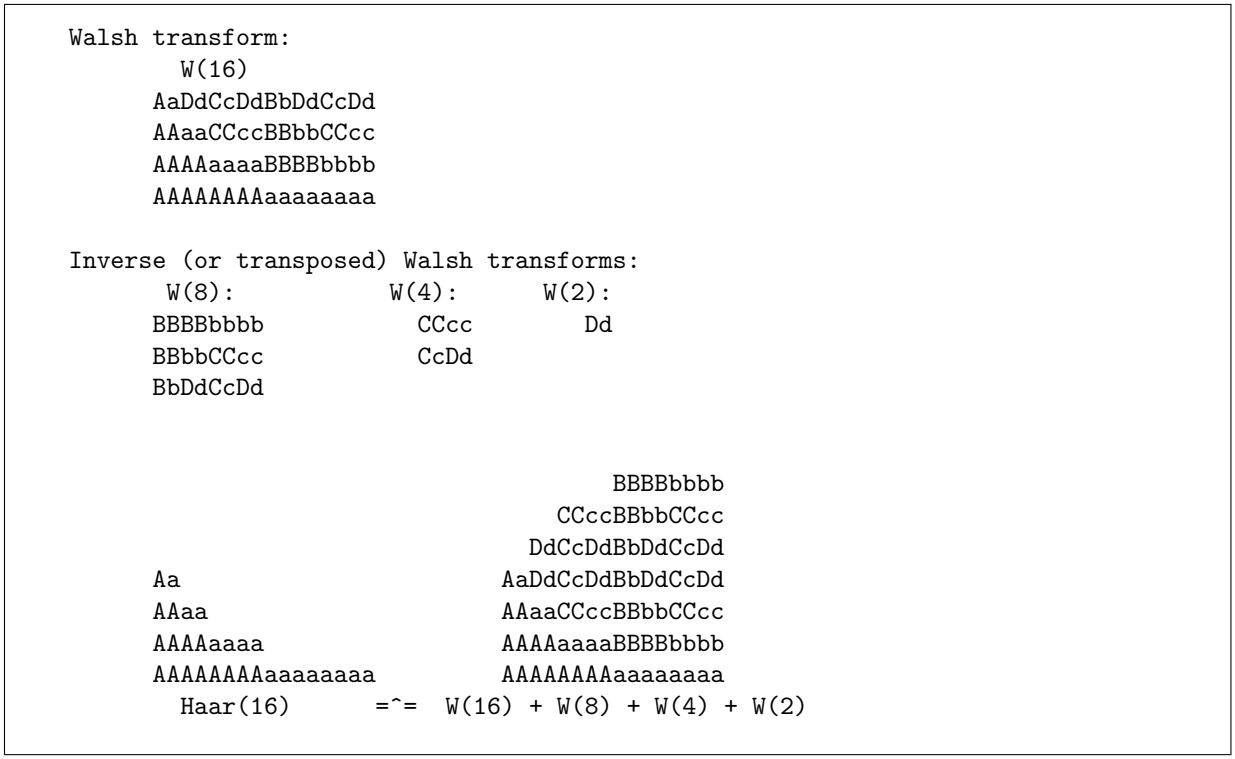

Figure 23.6-D: Symbolic description of how to build a Haar transform from Walsh transforms.

```
3       {
4           ulong k = 1UL << ldk;
5           walsh_pal(f+k, ldk);
6       }
7       walsh_pal(f, ldn);
```

and its inverse as

```
1       // algorithm HW2I:  inverse_transposed_haar_nn(f, ldn); =^=
2       walsh_pal(f, ldn);  // =^= revbin_permute(f, n); walsh_wak(f, ldn);
3       for (ulong ldk=1; ldk<ldn; ++ldk)
4       {
5           ulong k = 1UL << ldk;
6           walsh_pal(f+k, ldk);
7       }
```

The symbolic scheme is given in figure 23.6-D.

23.7 Nonstandard splitting schemes *

All radix-2 transforms recursively split the length of the array into halves. The size of the transforms is limited to powers of two. In a recursive implementation we use the fact that $2^k = 2^{k-1} + 2^{k-1}$. With $N_k := 2^k$ we have $N_0 = 1$, and $N_k = N_{k-1} + N_{k-1}$. We use different recursive schemes to derive nonstandard variants of the Haar and Walsh transforms.

23.7.1 Fibonacci-Haar and Fibonacci-Walsh transform

One can use the Fibonacci numbers $F_n = F_{n-1} + F_{n-1}$ (where $F_0 = 0$ and $F_1 = 1$) to construct a *Fibonacci-Haar transform* as follows [FXT: haar/fib-haar.h]:

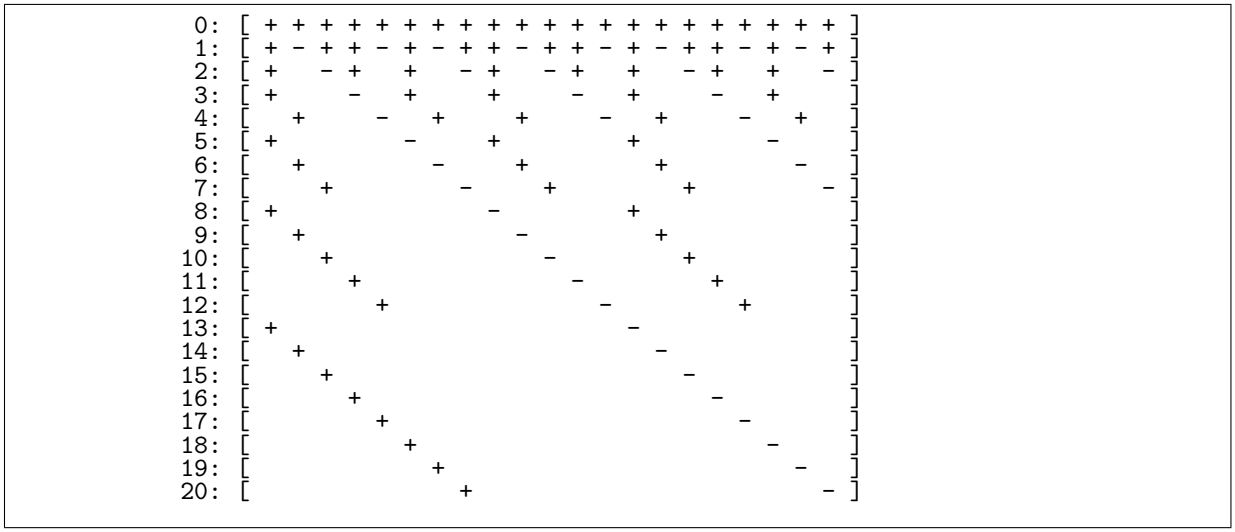

Figure 23.7-A: Basis functions for the non-normalized Fibonacci-Haar transform. Only the signs of the nonzero entries are shown. At the blank entries the functions are zero.

```
1   inline void fibonacci_haar(double *a, ulong f0, ulong f1)
2   // In-place Fibonacci-Haar transform of a[0,...,f0-1].
3   // f0 must be a Fibonacci number, f1 the next smaller Fibonacci number.
4   {
5       if ( f0 < 2 )  return;
6
7       ulong f2 = f0 - f1;
8       for (ulong j=0,k=f1;  j<f2;  ++j,++k)
9       {
10          double u = a[j],  v = a[k];
11          a[j] = (u+v) * SQRT1_2;
12          a[k] = (u-v) * SQRT1_2;
13      }
14      fibonacci_haar(a, f1, f2);
15  }
```

A non-normalized version is obtained by omitting the multiplications with $1/\sqrt{2}$ (=`SQRT1_2`). The basis functions for the non-normalized transform with length-21 ($= F_8$) are shown in figure 23.7-A (compare to figure 23.5-A on page 492). The second row corresponds to the *rabbit sequence* described in section 36.11 on page 740. Figure 23.7-A was created with the program [FXT: fft/fib-haar-demo.cc].

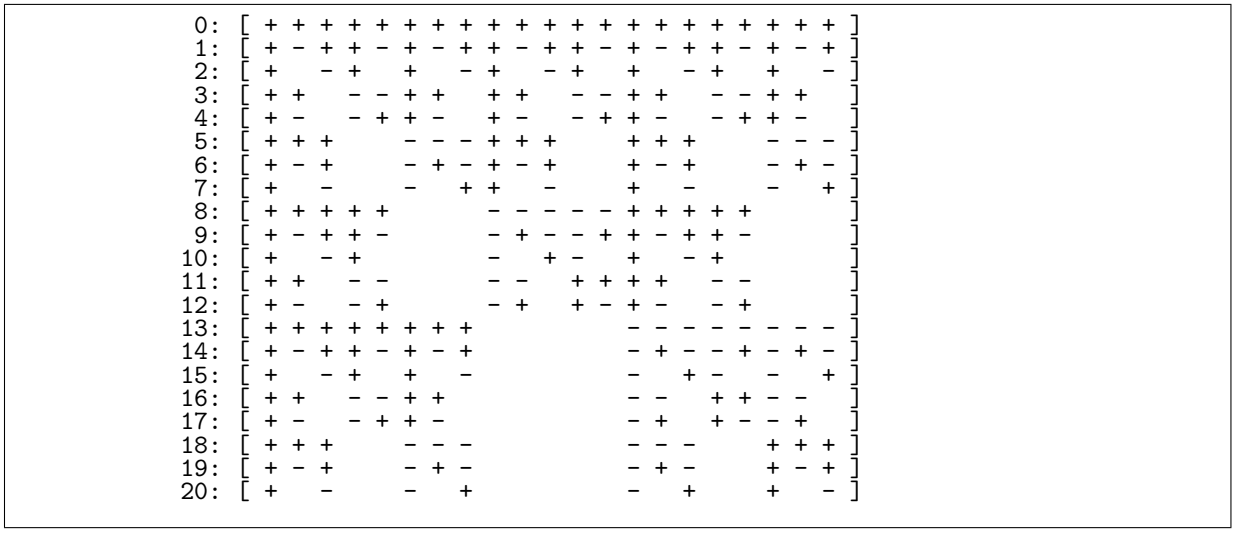

Figure 23.7-B: Basis functions for the non-normalized Fibonacci-Walsh transform.

A *Fibonacci-Walsh transform* can be obtained by adding one line in the recursive implementation of the Fibonacci-Haar transform [FXT: walsh/fib-walsh.h]:

```
inline void fibonacci_walsh(double *a, ulong f0, ulong f1)
// In-place Fibonacci-Walsh transform of a[0,...,f0-1].
// f0 must be a Fibonacci number, f1 the next smaller Fibonacci number.
{
    if ( f0 < 2 )  return;

    ulong f2 = f0 - f1;
    for (ulong j=0,k=f1;  j<f2;  ++j,++k)
    {
        double u = a[j],  v = a[k];
        a[j] = (u+v) * SQRT1_2;
        a[k] = (u-v) * SQRT1_2;
    }
    fibonacci_walsh(a, f1, f2);
    fibonacci_walsh(a+f1, f2, f1-f2);  // <--= omit line to obtain Haar transform
}
```

The basis functions for the length 21 transform are shown in figure 23.7-B which was created with the program [FXT: fft/fib-walsh-demo.cc].

The given routines can be optimized by inserting short-length transforms as recursion end.

One can obtain Haar-like and Walsh-like transforms for any linear recursive sequence that is increasing. A construction for recurrences $N_k = N_{k-1} + N_{k-1-p}$ is considered in [112].

23.7.2 Mersenne-Haar and Mersenne-Walsh transform

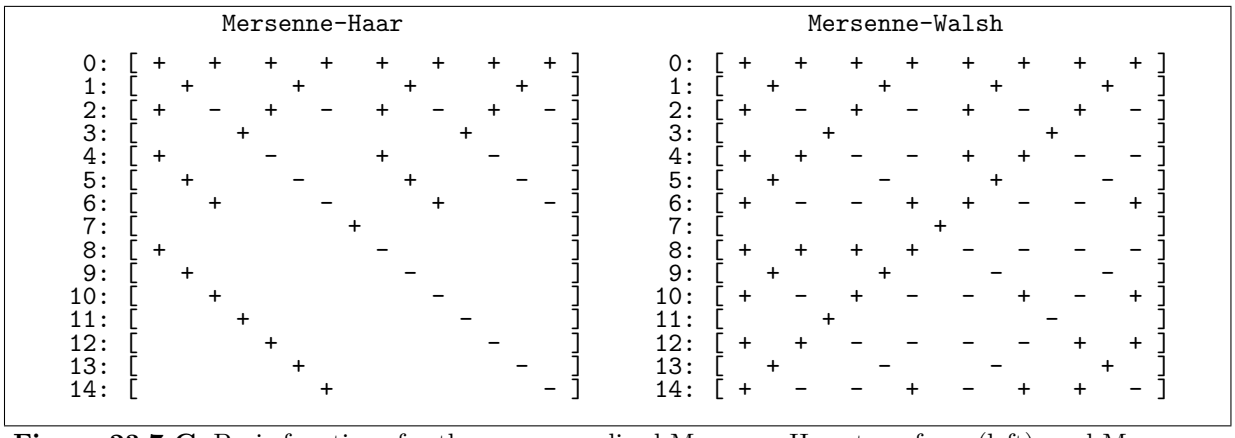

Figure 23.7-C: Basis functions for the non-normalized Mersenne-Haar transform (left), and Mersenne-Walsh transform (right). Only the signs of the nonzero entries are shown. At the blank entries the functions are zero.

For the Mersenne numbers $M_k = 2^k - 1$ we have the recursion $M_k = 2 \cdot M_{k-1} + 1$. This can be used to obtain a *Mersenne-Walsh transform* [FXT: walsh/mers-walsh.h]:

```
inline void mersenne_walsh(double *a, ulong f0)
// In-place Mersenne-Walsh transform of a[0,...,f0-1].
// f0 must be a Mersenne number.
// Self-inverse.
{
    if ( f0 < 2 )  return;
    ulong f1 = f0 >> 1;  // next smaller Mersenne number

    for (ulong j=0,k=f1+1;  j<f1;  ++j,++k)
    {
        double u = a[j],  v = a[k];
        a[j] = (u+v) * SQRT1_2;
        a[k] = (u-v) * SQRT1_2;
    }
```

```
15      mersenne_walsh(a, f1);
16      mersenne_walsh(a+f1+1, f1);  // <--= omit line to obtain Mersenne-Haar transform
17  }
```

Figure 23.7-C (right) gives the basis functions for the non-normalized Mersenne-Walsh transform. The Mersenne-Haar transform is obtained by deleting one line as indicated. The implementation is given in [FXT: haar/mers-haar.h], the basis functions of the non-normalized version are shown at the left of figure 23.7-C. The figure was created with the programs [FXT: fft/mers-walsh-demo.cc] and [FXT: fft/mers-haar-demo.cc]. Note that both transforms leave the central element unchanged.

Chapter 24

The Hartley transform

The Hartley transform is a trigonometric transform whose practical importance comes from the fact that it maps real data to real data. While the fast algorithms for radix-2 can be found without great difficulty the higher radix algorithms are not obvious. Therefore it is appropriate to describe the Hartley transform in terms of the Fourier transform. A method for the conversion of FFT algorithms to fast Hartley transform (FHT) algorithms is given.

Routines for the conversion of Hartley transforms to and from Fourier transforms are described. Convolution routines based on the FHT are given for complex and real valued data. An efficient procedure for the computation of the negacyclic convolution is described.

24.1 Definition and symmetries

The *discrete Hartley transform* (HT) of a length-n sequence a is defined as

$$c = \mathcal{H}[a] \tag{24.1-1a}$$

$$c_k := \frac{1}{\sqrt{n}} \sum_{x=0}^{n-1} a_x \left(\cos \frac{2\pi k x}{n} + \sin \frac{2\pi k x}{n} \right) \tag{24.1-1b}$$

That is, almost like the Fourier transform but with '$\cos + \sin$' instead of '$\cos + i \cdot \sin$'. The (continuous version of the) Hartley transform is treated in [142].

The Hartley transform of a purely real sequence is purely real:

$$\mathcal{H}[a] \in \mathbb{R} \qquad \text{for} \quad a \in \mathbb{R} \tag{24.1-2}$$

It also is its own inverse:

$$\mathcal{H}[\mathcal{H}[a]] = a \tag{24.1-3}$$

Symmetry is conserved, like for the Fourier transform: the Hartley transform of a symmetric, antisymmetric sequence is symmetric, antisymmetric, respectively. Using the notation from section 20.7 on page 414 one has

$$\mathcal{H}[a_S] = +\overline{\mathcal{H}[a_S]} = +\mathcal{H}[\overline{a_S}] \tag{24.1-4a}$$

$$\mathcal{H}[a_A] = -\overline{\mathcal{H}[a_A]} = -\mathcal{H}[\overline{a_A}] \tag{24.1-4b}$$

An algorithm for the fast ($n \log(n)$-) computation of the Hartley transform is called a *fast Hartley transform* (FHT).

24.2 Radix-2 FHT algorithms

24.2.1 Decimation in time (DIT) FHT

Notation: For a length-n sequence a of let $\mathcal{X}^{1/2}a$ denote the sequence with elements $a_x \cos \pi\, x/n + \overline{a}_x \sin \pi\, x/n$. The operator $\mathcal{X}^{1/2}$ is the equivalent to the operator $\mathcal{S}^{1/2}$ of the Fourier transform algorithms. We use the notation $(even)$ and (odd) as introduced on page 396. The radix-2 decimation in time (DIT) step for the FHT:

$$\mathcal{H}\,[a]^{(left)} \quad \stackrel{n/2}{=} \quad \mathcal{H}\left[a^{(even)}\right] + \mathcal{X}^{1/2}\mathcal{H}\left[a^{(odd)}\right] \tag{24.2-1a}$$

$$\mathcal{H}\,[a]^{(right)} \quad \stackrel{n/2}{=} \quad \mathcal{H}\left[a^{(even)}\right] - \mathcal{X}^{1/2}\mathcal{H}\left[a^{(odd)}\right] \tag{24.2-1b}$$

This is the equivalent to relations 20.3-3a and 20.3-3b on page 397.

Pseudo code for a recursive radix-2 DIT FHT (C++ version in [FXT: fht/recfht2.cc]):

```
1   procedure rec_fht_dit2(a[], n, x[])
2   // real a[0..n-1] input
3   // real x[0..n-1] result
4   {
5       real b[0..n/2-1], c[0..n/2-1]   // workspace
6       real s[0..n/2-1], t[0..n/2-1]   // workspace
7
8       if n == 1 then
9       {
10          x[0] := a[0]
11          return
12      }
13
14      nh := n/2;
15
16      for k:=0 to nh-1
17      {
18          s[k] := a[2*k]    // even indexed elements
19          t[k] := a[2*k+1]  // odd  indexed elements
20      }
21
22      rec_fht_dit2(s[], nh, b[])
23      rec_fht_dit2(t[], nh, c[])
24
25      hartley_shift(c[], nh, 1/2)
26
27      for k:=0 to nh-1
28      {
29          x[k]    := b[k] + c[k];
30          x[k+nh] := b[k] - c[k];
31      }
32  }
```

The result is returned in the array in `x[]`. The procedure `hartley_shift()` implements the operator $\mathcal{X}^{1/2}$, it replaces element c_k of the input sequence c by $c_k\, cos(\pi\, k/n) + c_{n-k}\, sin(\pi\, k/n)$. As pseudo code:

```
1   procedure hartley_shift_05(c[], n)
2   // real c[0..n-1] input, result
3   {
4       nh := n/2
5       j := n-1
6       for k:=1 to nh-1
7       {
8           c := cos( PI*k/n )
9           s := sin( PI*k/n )
10
11          {c[k], c[j]} := {c[k]*c+c[j]*s, c[k]*s-c[j]*c}
12
13          j := j-1
14      }
15  }
```

C++ implementations are given in [FXT: fht/hartleyshift.h]. A version that exploits the symmetry of the trigonometric factors is

```
#define Tdouble long double
#define Sin     sinl

template <typename Type>
inline void hartley_shift_05_v2rec(Type *f, ulong n)
{
    const ulong nh = n/2;
    if ( n>=4 )
    {
        ulong i0=nh/2, j0=3*i0;
        Type fi = f[i0],  fj = f[j0];
        double cs = SQRT1_2;
        f[i0] = (fi + fj) * cs;
        f[j0] = (fi - fj) * cs;

        if ( n>=8 )
        {
            const Tdouble phi0 = PI/n;
            Tdouble be = Sin(phi0),  al = Sin(0.5*phi0);  al *= (2.0*al);
            Tdouble s = 0.0,  c = 1.0;
            for (ulong i=1, j=n-1, k=nh-1, l=nh+1;  i<k;  ++i, --j, --k, ++l)
            {
                { Tdouble tt = c;  c -= (al*tt+be*s);  s -= (al*s-be*tt); }

                fi = f[i];
                fj = f[j];
                f[i] = fi * (double)c + fj * (double)s;
                f[j] = fi * (double)s - fj * (double)c;

                fi = f[k];
                fj = f[l];
                f[k] = fi * (double)s + fj * (double)c;
                f[l] = fi * (double)c - fj * (double)s;
            }
        }
    }
}

#undef Tdouble
#undef Sin
```

Pseudo code for a non-recursive radix-2 DIT FHT:

```
procedure fht_depth_first_dit2(a[], ldn)
// real a[0..n-1] input,result
{
    n := 2**ldn  // length of a[] is a power of 2

    revbin_permute(a[], n)

    for ldm:=1 to ldn
    {
        m  := 2**ldm
        mh := m/2
        m4 := m/4

        for r:=0 to n-m step m
        {
            for j:=1 to m4-1  // hartley_shift(a+r+mh,mh,1/2)
            {
                k := mh - j

                u := a[r+mh+j]
                v := a[r+mh+k]

                c := cos(j*PI/mh)
                s := sin(j*PI/mh)

                {u, v} := {u*c+v*s, u*s-v*c}

                a[r+mh+j] := u
                a[r+mh+k] := v
            }

            for j:=0 to mh-1
            {
                u := a[r+j]
                v := a[r+j+mh]

```

```
37                  a[r+j]    := u + v
38                  a[r+j+mh] := u - v
39              }
40          }
41      }
42  }
```

The derivation of the 'usual' DIT2 FHT algorithm starts by combining the Hartley-shift with the sum/diff-operations [FXT: fht/fhtdit2.cc]:

```
1   void fht_depth_first_dit2(double *f, ulong ldn)
2   {
3       const ulong n = 1UL<<ldn;
4
5       revbin_permute(f, n);
6
7       for (ulong ldm=1; ldm<=ldn; ++ldm)
8       {
9           const ulong m = (1UL<<ldm);
10          const ulong mh = (m>>1);
11          const ulong m4 = (mh>>1);
12          const double phi0 = M_PI/mh;
13
14          for (ulong r=0; r<n; r+=m)
15          {
16              {  // j == 0:
17                  ulong t1 = r;
18                  ulong t2 = t1 + mh;
19                  sumdiff(f[t1], f[t2]);
20              }
21
22              if ( m4 )
23              {
24                  ulong t1 = r + m4;
25                  ulong t2 = t1 + mh;
26                  sumdiff(f[t1], f[t2]);
27              }
28
29              for (ulong j=1, k=mh-1; j<k; ++j,--k)
30              {
31                  double s, c;
32                  SinCos(phi0*j, &s, &c);
33
34                  ulong tj = r + mh + j;
35                  ulong tk = r + mh + k;
36                  double fj = f[tj];
37                  double fk = f[tk];
38                  f[tj] = fj * c + fk * s;
39                  f[tk] = fj * s - fk * c;
40
41                  ulong t1 = r + j;
42                  ulong t2 = tj; // == t1 + mh;
43                  sumdiff(f[t1], f[t2]);
44
45                  t1 = r + k;
46                  t2 = tk; // == t1 + mh;
47                  sumdiff(f[t1], f[t2]);
48              }
49          }
50      }
51  }
```

Finally, as with the FFT equivalent (see page 399), the number of trigonometric computations can be reduced by swapping the innermost loops [FXT: fht/fhtdit2.cc]:

```
1   void fht_dit2(double *f, ulong ldn)
2   // Radix-2 decimation in time (DIT) FHT.
3   {
4       const ulong n = 1UL<<ldn;
5
6       revbin_permute(f, n);
7
8       for (ulong ldm=1; ldm<=ldn; ++ldm)
9       {
10          const ulong m = (1UL<<ldm);
11          const ulong mh = (m>>1);
12          const ulong m4 = (mh>>1);
```

```
13          const double phi0 = M_PI/mh;
14
15          for (ulong r=0; r<n; r+=m)
16          {
17              {  // j == 0:
18                  ulong t1 = r;
19                  ulong t2 = t1 + mh;
20                  sumdiff(f[t1], f[t2]);
21              }
22
23              if ( m4 )
24              {
25                  ulong t1 = r + m4;
26                  ulong t2 = t1 + mh;
27                  sumdiff(f[t1], f[t2]);
28              }
29          }
30
31          for (ulong j=1, k=mh-1; j<k; ++j,--k)
32          {
33              double s, c;
34              SinCos(phi0*j, &s, &c);
35
36              for (ulong r=0; r<n; r+=m)
37              {
38                  ulong tj = r + mh + j;
39                  ulong tk = r + mh + k;
40                  double fj = f[tj];
41                  double fk = f[tk];
42                  f[tj] = fj * c + fk * s;
43                  f[tk] = fj * s - fk * c;
44
45                  ulong t1 = r + j;
46                  ulong t2 = tj; // == t1 + mh;
47                  sumdiff(f[t1], f[t2]);
48
49                  t1 = r + k;
50                  t2 = tk; // == t1 + mh;
51                  sumdiff(f[t1], f[t2]);
52              }
53          }
54      }
55  }
```

24.2.2 Decimation in frequency (DIF) FHT

The radix-2 decimation in frequency step for the FHT is (compare to relations 20.3-6a and 20.3-6b on page 400):

$$\mathcal{H}\,[a]^{(even)} \quad \stackrel{n/2}{=} \quad \mathcal{H}\left[a^{(left)} + a^{(right)}\right] \tag{24.2-2a}$$

$$\mathcal{H}\,[a]^{(odd)} \quad \stackrel{n/2}{=} \quad \mathcal{H}\left[\mathcal{X}^{1/2}\left(a^{(left)} - a^{(right)}\right)\right] \tag{24.2-2b}$$

Pseudo code for a recursive radix-2 DIF FHT (the C++ equivalent is given in [FXT: fht/recfht2.cc]):

```
1   procedure rec_fht_dif2(a[], n, x[])
2   // real a[0..n-1] input
3   // real x[0..n-1] result
4   {
5       real b[0..n/2-1], c[0..n/2-1]   // workspace
6       real s[0..n/2-1], t[0..n/2-1]   // workspace
7
8       if n == 1 then
9       {
10          x[0] := a[0]
11          return
12      }
13
14      nh := n/2;
15
16      for k:=0 to nh-1
```

```
17      {
18          s[k] := a[k]     // 'left'  elements
19          t[k] := a[k+nh]  // 'right' elements
20      }
21
22      for k:=0 to nh-1
23      {
24          {s[k], t[k]} := {s[k]+t[k], s[k]-t[k]}
25      }
26
27      hartley_shift(t[], nh, 1/2)
28
29      rec_fht_dif2(s[], nh, b[])
30      rec_fht_dif2(t[], nh, c[])
31
32      j := 0
33      for k:=0 to nh-1
34      {
35          x[j]   := b[k]
36          x[j+1] := c[k]
37          j := j+2
38      }
39  }
```

Pseudo code for a non-recursive radix-2 DIF FHT (C++ version in [FXT: fht/fhtdif2.cc]):

```
1   procedure fht_depth_first_dif2(a[], ldn)
2   // real a[0..n-1] input,result
3   {
4       n := 2**ldn  // length of a[] is a power of 2
5
6       for ldm:=ldn to 1 step -1
7       {
8           m  := 2**ldm
9           mh := m/2
10          m4 := m/4
11
12          for r:=0 to n-m step m
13          {
14              for j:=0 to mh-1
15              {
16                  u := a[r+j]
17                  v := a[r+j+mh]
18
19                  a[r+j]    := u + v
20                  a[r+j+mh] := u - v
21              }
22
23              for j:=1 to m4-1
24              {
25                  k := mh - j
26
27                  u := a[r+mh+j]
28                  v := a[r+mh+k]
29
30                  c := cos(j*PI/mh)
31                  s := sin(j*PI/mh)
32
33                  {u, v} := {u*c+v*s, u*s-v*c}
34
35                  a[r+mh+j] := u
36                  a[r+mh+k] := v
37              }
38          }
39      }
40
41      revbin_permute(a[], n)
42  }
```

The 'usual' DIF2 FHT algorithm then is again obtained by swapping the inner loops, a C++ implementation is [FXT: `fht_dif2()` in fht/fhtdif2.cc]:

```
1   void fht_dif2(double *f, ulong ldn)
2   // Radix-2 decimation in frequency (DIF) FHT
3   {
4       const ulong n = (1UL<<ldn);
5       for (ulong ldm=ldn; ldm>=1; --ldm)
6       {
7           const ulong m = (1UL<<ldm);
```

```
 8          const ulong mh = (m>>1);
 9          const ulong m4 = (mh>>1);
10          const double phi0 = M_PI/mh;
11
12          for (ulong r=0; r<n; r+=m)
13          {
14              {  // j == 0:
15                  ulong t1 = r;
16                  ulong t2 = t1 + mh;
17                  sumdiff(f[t1], f[t2]);
18              }
19
20              if ( m4 )
21              {
22                  ulong t1 = r + m4;
23                  ulong t2 = t1 + mh;
24                  sumdiff(f[t1], f[t2]);
25              }
26          }
27
28          for (ulong j=1, k=mh-1; j<k; ++j,--k)
29          {
30              double s, c;
31              SinCos(phi0*j, &s, &c);
32
33              for (ulong r=0; r<n; r+=m)
34              {
35                  ulong tj = r + mh + j;
36                  ulong tk = r + mh + k;
37
38                  ulong t1 = r + j;
39                  ulong t2 = tj; // == t1 + mh;
40                  sumdiff(f[t1], f[t2]);
41
42                  t1 = r + k;
43                  t2 = tk; // == t1 + mh;
44                  sumdiff(f[t1], f[t2]);
45
46                  double fj = f[tj];
47                  double fk = f[tk];
48                  f[tj] = fj * c + fk * s;
49                  f[tk] = fj * s - fk * c;
50              }
51          }
52      }
53
54      revbin_permute(f, n);
55  }
```

24.3 Complex FT by HT

The relations between the HT and the FT can be read off directly from their definitions and their symmetry relations. Let σ be the sign of the exponent in the FT, then the HT of a complex sequence $d \in \mathbb{C}$ is

$$\mathcal{F}\big[\,d\,\big] \;=\; \frac{1}{2}\left(\mathcal{H}\,[d] + \overline{\mathcal{H}\,[d]} + \sigma\, i\,\left(\mathcal{H}\,[d] - \overline{\mathcal{H}\,[d]}\right)\right) \tag{24.3-1}$$

Written out for the real and imaginary part of $d = a + i\,b$ $(a, b \in \mathbb{R})$:

$$\mathfrak{Re}\,\mathcal{F}\big[\,a + i\,b\,\big] \;=\; \frac{1}{2}\left(\mathcal{H}\,[a] + \overline{\mathcal{H}\,[a]} - \sigma\,\left(\mathcal{H}\,[b] - \overline{\mathcal{H}\,[b]}\right)\right) \tag{24.3-2a}$$

$$\mathfrak{Im}\,\mathcal{F}\big[\,a + i\,b\,\big] \;=\; \frac{1}{2}\left(\mathcal{H}\,[b] + \overline{\mathcal{H}\,[b]} + \sigma\,\left(\mathcal{H}\,[a] - \overline{\mathcal{H}\,[a]}\right)\right) \tag{24.3-2b}$$

Using the symmetry relations 24.1-4a and 24.1-4b on page 501 one can recast the relations as

$$\mathfrak{Re}\,\mathcal{F}\big[\,a + i\,b\,\big] \;=\; \frac{1}{2}\,\mathcal{H}\,[a_S - \sigma\,b_A] \tag{24.3-3a}$$

$$\mathfrak{Im}\,\mathcal{F}\big[\,a + i\,b\,\big] \;=\; \frac{1}{2}\,\mathcal{H}\,[b_S + \sigma\,a_A] \tag{24.3-3b}$$

Both formulations lead to the very same conversion procedure. The following pseudo code is for a complex FT by HT conversion:

```
fht_fft_conversion(a[], b[], n, is)
// preprocessing to use two length-n FHTs
// to compute a length-n complex FFT
// or
// postprocessing to use two length-n FHTs
// to compute a length-n complex FFT
//
// Self-inverse.
{
    for k:=1 to n/2-1
    {
        t := n-k

        as := a[k] + a[t]
        aa := a[k] - a[t]

        bs := b[k] + b[t]
        ba := b[k] - b[t]

        aa := is * aa
        ba := is * ba

        a[k] := 1/2 * (as - ba)
        a[t] := 1/2 * (as + ba)

        b[k] := 1/2 * (bs + aa)
        b[t] := 1/2 * (bs - aa)
    }
}
```

The C++ implementations are given in [FXT: fft/fhtfft.cc] for type `double` and [FXT: fft/fhtcfft.cc] for type `complex`. Now we have two options to compute a complex FT by two HTs. Version 1 does the FHTs first:

```
fft_by_fht1(a[], b[], n, is)
// real a[0..n-1] input,result (real part)
// real b[0..n-1] input,result (imaginary part)
{
    fht(a[], n)
    fht(b[], n)
    fht_fft_conversion(a[], b[], n, is)
}
```

Version 2 does the FHTs at the end of the routine:

```
fft_by_fht2(a[], b[], n, is)
// real a[0..n-1] input,result (real part)
// real b[0..n-1] input,result (imaginary part)
{
    fht_fft_conversion(a[], b[], n, is)
    fht(a[], n)
    fht(b[], n)
}
```

Note that the real and imaginary parts of the FT are computed independently by this procedure. This can be very advantageous when the real and imaginary part of complex data lies in separate arrays. The C++ version is given in [FXT: fft/fhtfft.cc].

24.4 Complex FT by complex HT and vice versa

A complex valued HT is simply two HTs (one of the real, one of the imaginary part). So we can use either version from section 24.3 and there is nothing new. Really? If one has a type `complex` version of both the conversion and the FHT routine then the complex FFT can be computed as either

```
fft_by_fht1(c[], n, is)
// complex c[0..n-1] input,result
{
    fht(c[], n)
    fht_fft_conversion(c[], n, is)
}
```

or the same with swapped statements.

This may not make you scream but here is the message: it makes sense to do so. One saves half of the trigonometric computations and book keeping. It is pretty easy to derive a complex FHT from the real version and with a well optimized FHT you get an even better optimized FFT. C++ implementations of complex FHTs are given in [FXT: fht/cfhtdif.cc] (DIF algorithm), [FXT: fht/cfhtdit.cc] (DIT algorithm), and, for zero padded data, [FXT: fht/cfht0.cc].

The other way round: computation of a complex FHT using FFTs. Let T be the operator corresponding to the `fht_fft_conversion`. The operator is its own inverse: $T = T^{-1}$. We have seen that

$$\mathcal{F} = \mathcal{H} \cdot T \qquad \text{and} \qquad \mathcal{F} = T \cdot \mathcal{H} \tag{24.4-1}$$

Thereby (multiply the relations with T and use $T \cdot T = 1$):

$$\mathcal{H} = T \cdot \mathcal{F} \qquad \text{and} \qquad \mathcal{H} = \mathcal{F} \cdot T \tag{24.4-2}$$

Hence we have either

```
1  fht_by_fft(c[], n, is)
2  // complex c[0..n-1] input,result
3  {
4      fft(c[], n)
5      fht_fft_conversion(c[], n, is)
6  }
```

or the same thing with swapped lines [FXT: fft/fhtcfft.cc]. The same ideas also work for separate real and imaginary parts but in that case one should rather use separate FHTs for the two arrays.

24.5 Real FT by HT and vice versa

To express the real and imaginary part of a Fourier transform of a purely real sequence $a \in \mathbb{R}$ by its Hartley transform use relations 24.3-2a and 24.3-2b on page 507 and set $b = 0$:

$$\mathfrak{Re}\,\mathcal{F}\big[\,a\,\big] \;=\; \frac{1}{2}\,(\mathcal{H}\,[a] + \overline{\mathcal{H}\,[a]}) \tag{24.5-1a}$$

$$\mathfrak{Im}\,\mathcal{F}\big[\,a\,\big] \;=\; \sigma\,\frac{1}{2}\,(\mathcal{H}\,[a] - \overline{\mathcal{H}\,[a]}) \tag{24.5-1b}$$

A C++ implementation is [FXT: `fht_real_complex_fft()` in realfft/realfftbyfht.cc]:

```
1  template <typename Type>
2  static inline void sumdiff05(Type &a, Type &b)
3  // {a, b}  <--| {0.5*(a+b),  0.5*(a-b)}
4  { Type t=(a-b)*0.5; a+=b; a*=0.5; b=t; }
5
6  template <typename Type>
7  static inline void sumdiff05_r(Type &a, Type &b)
8  // {a, b}  <--| {0.5*(a+b), 0.5*(b-a)}
9  { Type t=(b-a)*0.5; a+=b; a*=0.5; b=t; }
10
11 void
12 fht_real_complex_fft(double *f, ulong ldn, int is/*=+1*/)
13 {
14     fht(f, ldn);
15
16     const ulong n = (1UL<<ldn);
17
18     if ( is>0 )  for (ulong i=1,j=n-1; i<j; i++,j--)  sumdiff05(f[i], f[j]);
19     else         for (ulong i=1,j=n-1; i<j; i++,j--)  sumdiff05_r(f[i], f[j]);
20 }
```

At the end of the procedure the ordering of the output data $c = \mathcal{F}\big[\, a \,\big] \in \mathbb{C}$ is

$$\begin{aligned}
\texttt{a[0]} &= \mathfrak{Re}\, c_0 \\
\texttt{a[1]} &= \mathfrak{Re}\, c_1 \\
\texttt{a[2]} &= \mathfrak{Re}\, c_2 \\
&\dots \\
\texttt{a[n/2]} &= \mathfrak{Re}\, c_{n/2} \\
\texttt{a[n/2+1]} &= \mathfrak{Im}\, c_{n/2-1} \\
\texttt{a[n/2+2]} &= \mathfrak{Im}\, c_{n/2-2} \\
\texttt{a[n/2+3]} &= \mathfrak{Im}\, c_{n/2-3} \\
&\dots \\
\texttt{a[n-1]} &= \mathfrak{Im}\, c_1
\end{aligned} \tag{24.5-2}$$

The inverse procedure is given in [FXT: realfft/realfftbyfht.cc]:

```
1   void
2   fht_complex_real_fft(double *f, ulong ldn, int is/*=+1*/)
3   {
4       const ulong n = (1UL<<ldn);
5
6       if ( is>0 )  for (ulong i=1,j=n-1; i<j; i++,j--)  sumdiff(f[i], f[j]);
7       else         for (ulong i=1,j=n-1; i<j; i++,j--)  diffsum(f[i], f[j]);
8
9       fht(f,ldn);
10  }
```

The function `sumdiff()` is defined in [FXT: aux0/sumdiff.h]:

```
1   template <typename Type>
2   static inline void sumdiff(Type &a, Type &b)
3   // {a, b}  <--| {a+b, a-b}
4   { Type t=a-b; a+=b; b=t; }
5
6   template <typename Type>
7   static inline void diffsum(Type &a, Type &b)
8   // {a, b}  <--| {a-b, a+b}
9   { Type t=a-b; b+=a; a=t; }
```

The input has to be ordered as given above (relations 24.5-2). The sign of the transform (`is`) has to be the same as with the forward version.

Computation of a (real-valued) FHT using a real-valued FFT proceeds similar as for complex versions. Let T_{r2c} be the operator corresponding to the post-processing in `real_complex_fft_by_fht()`, and T_{c2r} correspond to the preprocessing in `complex_real_fft_by_fht()`. That is

$$\mathcal{F}_{c2r} = \mathcal{H} \cdot T_{c2r} \qquad \text{and} \qquad \mathcal{F}_{r2c} = T_{r2c} \cdot \mathcal{H} \tag{24.5-3}$$

The operators are mutually inverse: $T_{r2c} = T_{c2r}^{-1}$ and $T_{c2r} = T_{r2c}^{-1}$. Multiplying the relations and using $T_{r2c} \cdot T_{c2r} = T_{c2r} \cdot T_{r2c} = 1$ gives

$$\mathcal{H} = T_{c2r} \cdot \mathcal{F}_{r2c} \qquad \text{and} \qquad \mathcal{H} = \mathcal{F}_{c2r} \cdot T_{r2c} \tag{24.5-4}$$

The corresponding code should be obvious. Watch out for real-to-complex FFTs that use a different ordering of the output than given in relation 24.5-2.

24.6 Higher radix FHT algorithms

Higher radix FHT algorithms seem to get complicated due to the structure of the Hartley shift operator. In fact there is a straightforward way to turn any FFT decomposition into an FHT algorithm.

For the moment assume that we want to compute a complex HT, further assume we want to use a radix-r algorithm. At each step we have r short HTs and want to combine them to a longer HT but we do not know how this might be done. In section 24.3 on page 507 we learned how to turn a HT into an FT using the T-operator. And we have seen radix-r algorithms for the FFT. The crucial idea is to use the conversion operator T as a wrapper around the FFT-step that combines several short FTs into a longer one. Here is how to turn a radix-r FFT-step into an FHT-step, simply do the following:

1. first convert the r short HTs into FTs (use T on the subsequences)
2. then perform the radix-r the FFT step
3. finally convert the FT into a HT (use T on the sequence)

For efficient implementations one obviously wants to combine the computations.

To obtain real-valued FHTs note that the real and imaginary parts do not 'mix': one can use the identical algorithm with real input (and the corresponding data types). With a radix-r step the scheme always accesses $2r$ elements simultaneously. The symmetry of the trigonometric factors is thereby automatically exploited. Splitting steps for the radix-4 FHT and the split-radix FHT are given in [244].

24.7 Convolution via FHT

The convolution property of the Hartley transform can be stated as

$$\mathcal{H}\left[a \circledast b\right] \quad = \quad \frac{1}{2}\Big(\mathcal{H}\left[a\right]\mathcal{H}\left[b\right] - \overline{\mathcal{H}\left[a\right]}\,\overline{\mathcal{H}\left[b\right]} + \mathcal{H}\left[a\right]\overline{\mathcal{H}\left[b\right]} + \overline{\mathcal{H}\left[a\right]}\,\mathcal{H}\left[b\right]\Big) \tag{24.7-1}$$

or, with $c := \mathcal{H}\left[a\right]$ and $d := \mathcal{H}\left[b\right]$, written element-wise:

$$\mathcal{H}\left[a \circledast b\right]_k \quad = \quad \frac{1}{2}\Big(c_k\,d_k - \overline{c_k}\,\overline{d_k} + c_k\,\overline{d_k} + \overline{c_k}\,d_k\Big) \tag{24.7-2a}$$

$$= \quad \frac{1}{2}\Big(c_k\,(d_k + \overline{d_k}) + \overline{c_k}\,(d_k - \overline{d_k})\Big) \tag{24.7-2b}$$

The latter form reduces the number of multiplications. When turning the relation into an algorithm one has to keep in mind that both elements $y_k = \mathcal{H}\left[a \circledast b\right]$ and y_{-k} must be computed simultaneously.

24.7.1 Implementation as pseudo code

Pseudo code for the cyclic convolution of two real valued sequences `x[]` and `y[]` via the FHT. `n` must be even, the result is returned in `y[]`:

```
1   procedure fht_cyclic_convolution(x[], y[], n)
2   // real x[0..n-1] input, modified
3   // real y[0..n-1] result
4   {
5       // transform data:
6       fht(x[], n)
7       fht(y[], n)
8
9       // convolution in transformed domain:
10      j := n-1
11      for i:=1 to n/2-1
12      {
13          xi := x[i]
14          xj := x[j]
15
16          yp := y[i] + y[j]   // ==   y[j] + y[i]
17          ym := y[i] - y[j]   // == -(y[j] - y[i])
18
19          y[i] := (xi*yp + xj*ym)/2
20          y[j] := (xj*yp - xi*ym)/2
```

```
21
22          j := j-1
23      }
24      y[0] := x[0]*y[0]
25      if n>1 then  y[n/2] := x[n/2]*y[n/2]
26
27      // transform back:
28      fht(y[], n)
29
30      // normalize:
31      for i:=0 to n-1
32      {
33          y[i] := y[i] / n
34      }
35  }
```

It is assumed that the procedure `fht()` does no normalization. The C++ equivalent is given in [FXT: convolution/fhtcnvl.cc].

Equation 24.7-2a (slightly optimized) for the auto convolution is

$$\mathcal{H}\left[a \circledast a\right]_k = \frac{1}{2}\Big(c_k\left(c_k + \overline{c_k}\right) + \overline{c_k}\left(c_k - \overline{c_k}\right)\Big) \tag{24.7-3a}$$

$$= c_k\,\overline{c_k} + \frac{1}{2}\big(c_k^2 - \overline{c_k}^2\big) \tag{24.7-3b}$$

where $c = \mathcal{H}\left[a\right]$.

We give pseudo code for the cyclic auto convolution that uses a fast Hartley transform, `n` must be even:

```
1   procedure cyclic_self_convolution(x[], n)
2   // real x[0..n-1] input, result
3   {
4       // transform data:
5       fht(x[], n)
6
7       // convolution in transformed domain:
8       j := n-1
9       for i:=1 to n/2-1
10      {
11          ci := x[i]
12          cj := x[j]
13
14          t1 := ci*cj               // == cj*ci
15          t2 := 1/2*(ci*ci-cj*cj)   // == -1/2*(cj*cj-ci*ci)
16
17          x[i] := t1 + t2
18          x[j] := t1 - t2
19
20          j := j-1
21      }
22      x[0]   := x[0]*x[0]
23
24      if n>1 then  x[n/2] := x[n/2]*x[n/2]
25
26      // transform back:
27      fht(x[], n)
28
29      // normalize:
30      for i:=0 to n-1
31      {
32          x[i] := x[i] / n
33      }
34  }
```

For odd `n` replace the line

```
for i:=1 to n/2-1
```

by

```
for i:=1 to (n-1)/2
```

and omit the line

```
if n>1 then  x[n/2] := x[n/2]*x[n/2]
```

in both procedures above.

24.7.2 C++ implementations

The FHT based routine for the cyclic convolution of two real sequences is:

```
1  void fht_convolution(double * restrict f, double * restrict g, ulong ldn)
2  {
3      fht(f, ldn);
4      fht(g, ldn);
5      fht_convolution_core(f, g, ldn);
6      fht(g, ldn);
7  }
```

The equivalent to the element-wise multiplication is given in [FXT: convolution/fhtcnvlcore.cc]:

```
1  void
2  fht_convolution_core(const double * restrict f, double * restrict g, ulong ldn,
3                       double v/*=0.0*/)
4  // Auxiliary routine for the computation of convolutions
5  //   via Fast Hartley Transforms.
6  // ldn := base-2 logarithm of the array length.
7  // v!=0.0 chooses alternative normalization.
8  {
9      const ulong n = (1UL<<ldn);
10
11     if ( v==0.0 )  v = 1.0/n;
12
13     g[0]  *=  (v * f[0]);
14     const ulong  nh = n/2;
15
16     if ( nh>0 )
17     {
18         g[nh] *= (v * f[nh]);
19         v *= 0.5;
20         for (ulong i=1,j=n-1; i<j; i++,j--)  fht_mul(f[i], f[j], g[i], g[j], v);
21     }
22 }
```

where [FXT: convolution/fhtmulsqr.h]:

```
1  template <typename Type>
2  static inline  void
3  fht_mul(Type xi, Type xj, Type &yi, Type &yj, double v)
4  // yi <-- v*( 2*xi*xj + xi*xi - xj*xj )
5  // yj <-- v*( 2*xi*xj - xi*xi + xj*xj )
6  {
7      Type h1p = xi,  h1m = xj;
8      Type s1 = h1p + h1m,  d1 = h1p - h1m;
9      Type h2p = yi,  h2m = yj;
10     yi = (h2p * s1 + h2m * d1) * v;
11     yj = (h2m * s1 - h2p * d1) * v;
12 }
```

A C++ implementation of the FHT based self-convolution is given in [FXT: convolution/fhtcnvla.cc]. It uses the routine

```
1  void
2  fht_auto_convolution_core(double *f, ulong ldn,
3                            double v/*=0.0*/)
4  // v!=0.0 chooses alternative normalization
5  {
6      const ulong  n  = (1UL<<ldn);
7      if ( v==0.0 )  v = 1.0/n;
8      f[0] *= (v * f[0]);
9      if ( n>=2 )
10     {
11         const ulong  nh = n/2;
12         f[nh] *= (v * f[nh]);
13         v *= 0.5;
14         for (ulong i=1,j=n-1; i<nh; i++,j--)  fht_sqr(f[i], f[j], v);
15     }
16 }
```

where [FXT: convolution/fhtmulsqr.h]:

```
1  template <typename Type>
2  static inline  void
3  fht_sqr(Type &xi, Type &xj, double v)
```

```
4   // xi <-- v*( 2*xi*xj + xi*xi - xj*xj )
5   // xj <-- v*( 2*xi*xj - xi*xi + xj*xj )
6   {
7       Type a = xi,  b = xj;
8       Type s1 = (a + b) * (a - b);
9       a *= b;
10      a += a;
11      xi = (a+s1) * v;
12      xj = (a-s1) * v;
13  }
```

24.7.3 Avoiding the revbin permutations

The observation that the revbin permutations can be omitted with FFT based convolutions (see section 21.1.3 on page 429) applies again [FXT: convolution/fhtcnvlcore.cc]:

```
1   void
2   fht_convolution_revbin_permuted_core(const double * restrict f,
3                                        double * restrict g,
4                                        ulong ldn,
5                                        double v/*=0.0*/)
6   // Same as fht_convolution_core() but with data access in revbin order.
7   {
8       const ulong n = (1UL<<ldn);
9
10      if ( v==0.0 )  v = 1.0/n;
11
12      g[0] *= (v * f[0]);  // 0 == revbin(0)
13      if ( n>=2 )  g[1] *= (v * f[1]); // 1 == revbin(nh)
14
15      if ( n<4 )  return;
16
17      v *= 0.5;
18      const ulong nh = (n>>1);
19
20      ulong r=nh, rm=n-1; // nh == revbin(1),  n1-1 == revbin(n-1)
21      fht_mul(f[r], f[rm], g[r], g[rm], v);
22
23      ulong k=2, km=n-2;
24      while ( k<nh  )
25      {
26          // k even:
27          rm -= nh;
28          ulong tr = r;
29          r^=nh;  for (ulong m=(nh>>1); !((r^=m)&m); m>>=1)  {;}
30          fht_mul(f[r], f[rm], g[r], g[rm], v);
31          --km;
32          ++k;
33
34          // k odd:
35          rm += (tr-r);
36          r += nh;
37          fht_mul(f[r], f[rm], g[r], g[rm], v);
38          --km;
39          ++k;
40      }
41  }
```

The optimized version saving three revbin permutations is [FXT: convolution/fhtcnvl.cc]:

```
1   void fht_convolution(double * restrict f, double * restrict g, ulong ldn)
2   {
3       fht_dif_core(f, ldn);
4       fht_dif_core(g, ldn);
5       fht_convolution_revbin_permuted_core(f, g, ldn);
6       fht_dit_core(g, ldn);
7   }
```

24.8 Negacyclic convolution via FHT

Pseudo code for the computation of the negacyclic (auto-) convolution via FHT:

```
1   procedure negacyclic_self_convolution(x[], n)
2   // real x[0..n-1]  input, result
3   {
4       hartley_shift_05(x, n)    // preprocess
5       fht(x, n)    // transform data
6
7       // convolution in transformed domain:
8       j := n-1
9       for i:=0 to n/2-1  // here i starts from zero
10      {
11          a := x[i]
12          b := x[j]
13
14          x[i] := a*b+(a*a-b*b)/2
15          x[j] := a*b-(a*a-b*b)/2
16          j := j-1
17      }
18
19      fht(x, n)    // transform back
20      hartley_shift_05(x, n)    // postprocess
21  }
```

C++ implementations for the negacyclic convolution and self convolution are given in [FXT: convolution/fhtnegacnvl.cc]. The FHT-based negacyclic convolution is turns out to be extremely useful for the computation of weighted transforms, for example in the MFA-based convolution for real input, see section 21.4.1 on page 439.

24.9 Localized FHT algorithms

Localized routines for the FHT can be obtained by slight modifications of the corresponding algorithms for the Walsh transform described in section 22.6 on page 459. The decimation in time (DIT) version is [FXT: fht/fhtloc2.h]:

```
1   template <typename Type>
2   void fht_loc_dit2_core(Type *f, ulong ldn)
3   {
4       if ( ldn<=13 )  // sizeof(Type)*(2**thres) <= L1_CACHE_BYTES
5       {
6           fht_dit_core(f, ldn);
7           return;
8       }
9
10      // Recursion:
11      fht_dit_core_2(f+2);  // ldm==1
12      fht_dit_core_4(f+4);  // ldm==2
13      fht_dit_core_8(f+8);  // ldm==3
14      for (ulong ldm=4; ldm<ldn; ++ldm)  fht_loc_dit2_core(f+(1UL<<ldm), ldm);
15
16
17      for (ulong ldm=1; ldm<=ldn; ++ldm)
18      {
19          const ulong m = (1UL<<ldm);
20          const ulong mh = (m>>1);
21          hartley_shift_05(f+mh, mh);
22          for (ulong t1=0, t2=mh;  t1<mh;  ++t1, ++t2)  sumdiff(f[t1], f[t2]);
23      }
24  }
```

The routine `hartley_shift_05()` is described in 24.2.1 on page 502. One should choose a implementation that uses trigonometric recursion as this improves performance considerably.

The decimation in frequency (DIF) version is:

```
1   template <typename Type>
2   void fht_loc_dif2_core(Type *f, ulong ldn)
```

```
3   {
4       if ( ldn<=13 )  // sizeof(Type)*(2**thres) <= L1_CACHE_BYTES
5       {
6           fht_dif_core(f, ldn);
7           return;
8       }
9
10      for (ulong ldm=ldn; ldm>=1; --ldm)
11      {
12          const ulong m = (1UL<<ldm);
13          const ulong mh = (m>>1);
14          for (ulong t1=0, t2=mh;  t1<mh;  ++t1, ++t2)  sumdiff(f[t1], f[t2]);
15          hartley_shift_05(f+mh, mh);
16      }
17
18      // Recursion:
19      fht_dif_core_2(f+2);  // ldm==1
20      fht_dif_core_4(f+4);  // ldm==2
21      fht_dif_core_8(f+8);  // ldm==3
22      for (ulong ldm=4; ldm<ldn; ++ldm)  fht_loc_dif2_core(f+(1UL<<ldm), ldm);
23  }
```

The (generated) short-length transforms are given in the files [FXT: fht/shortfhtdifcore.h] and [FXT: fht/shortfhtditcore.h]. For example, the length-8 decimation in frequency routine is

```
1   template <typename Type>
2   inline void
3   fht_dif_core_8(Type *f)
4   {
5      Type g0, f0, f1, g1;
6      sumdiff(f[0], f[4], f0, g0);
7      sumdiff(f[2], f[6], f1, g1);
8      sumdiff(f0, f1);
9      sumdiff(g0, g1);
10     Type s1, c1, s2, c2;
11     sumdiff(f[1], f[5], s1, c1);
12     sumdiff(f[3], f[7], s2, c2);
13     sumdiff(s1, s2);
14     sumdiff(f0, s1, f[0], f[1]);
15     sumdiff(f1, s2, f[2], f[3]);
16     c1 *= SQRT2;
17     c2 *= SQRT2;
18     sumdiff(g0, c1, f[4], f[5]);
19     sumdiff(g1, c2, f[6], f[7]);
20  }
```

An additional revbin permutation is needed if the data is required in order. The FHT can be computed by either

```
fht_loc_dif2_core(f, ldn);
revbin_permute(f, 1UL<<ldn);
```

or

```
revbin_permute(f, 1UL<<ldn);
fht_loc_dit2_core(f, ldn);
```

Performance for large arrays is excellent: the convolutions based on the transforms [FXT: convolution/fhtloccnvl.cc]

```
1   void
2   loc_fht_convolution(double * restrict f, double * restrict g, ulong ldn)
3   {
4       fht_loc_dif2_core(f, ldn);
5       fht_loc_dif2_core(g, ldn);
6       fht_convolution_revbin_permuted_core(f, g, ldn);
7       fht_loc_dit2_core(g, ldn);
8   }
```

and [FXT: convolution/fhtloccnvla.cc]

```
1   void
2   loc_fht_auto_convolution(double *f, ulong ldn)
3   {
4       fht_loc_dif2_core(f, ldn);
5       fht_auto_convolution_revbin_permuted_core(f, ldn);
6       fht_loc_dit2_core(f, ldn);
```

```
7   }
```

gave a significant (more than 50 percent) speedup for the high precision multiplication routines (see section 27.3 on page 550) used in the hfloat library [20].

24.10 Two-dimensional FHTs

A two-dimensional FHT can be computed almost as easy as a two-dimensional FFT, only a trivial additional step is needed. Start with the row-column algorithm described in section 20.10.2 on page 423 [FXT: fht/twodimfht.cc]:

```
1   void
2   row_column_fht(double *f, ulong nr, ulong nc)
3   // FHT over rows and columns.
4   // nr := number of rows
5   // nc := number of columns
6   {
7       ulong n = nr * nc;
8
9       // fht over rows:
10      ulong ldc = ld(nc);
11      for (ulong k=0; k<n; k+=nc)  FHT(f+k, ldc);
12
13      // fht over columns:
14      double *w = new double[nr];
15      for (ulong k=0; k<nc; k++)  skip_fht(f+k, nr, nc, w);
16      delete [] w;
17  }
```

Note that no attempt has been made to make the routine cache friendly: the routine `skip_fht()` [FXT: fht/skipfht.cc] simply copies a column into the scratch array, does the FHT and copies the data back. This is not yet a two-dimensional FHT, the following post-processing must be made:

```
1   void
2   y_transform(double *f, ulong nr, ulong nc)
3   // Transforms row-column-FHT to 2-dimensional FHT.
4   // Self-inverse.
5   // nr := number of rows
6   // nc := number of columns
7   {
8       ulong rh = nr/2;
9       if ( nr&1 )  rh++;
10
11      ulong ch = nc/2;
12      if ( nc&1 )  ch++;
13
14      ulong n = nr*nc;
15      for (ulong tr=1, ctr=nc; tr<rh; tr++,ctr+=nc) // ctr=nc*tr
16      {
17          double *pa = f + ctr;
18          double *pb = pa + nc;
19          double *pc = f + n - ctr;
20          double *pd = pc + nc;
21          for (ulong tc=1; tc<ch; tc++)
22          {
23              pa++;
24              pb--;
25              pc++;
26              pd--;
27              double e = (*pa + *pd - *pb - *pc) * 0.5;
28              *pa -= e;
29              *pb += e;
30              *pc += e;
31              *pd -= e;
32          }
33      }
34  }
```

The canned routine is therefore

```
1   void
2   twodim_fht(double *f, ulong nr, ulong nc)
3   // Two dimensional fast Hartley transform (FHT)
```

```
4   // nr := number of rows
5   // nc := number of columns
6   {
7       row_column_fht(f, nr, nc);
8       y_transform(f, nr, nc);
9   }
```

24.11 Discrete cosine transform (DCT) by HT

The discrete cosine transform (DCT) with respect to the basis

$$u(k) = \nu(k) \cdot \cos\left(\frac{\pi\, k\,(i+1/2)}{n}\right) \quad \text{where} \tag{24.11-1}$$

$$\nu(k) = \begin{cases} 1 & \text{if} \quad k = 0 \\ \sqrt{2} & \text{else} \end{cases}$$

can be computed from the FHT using an auxiliary routine which is its own inverse. As pseudo code:

```
1   procedure cos_rot(x[], y[], n)
2   // Real x[0..n-1] input
3   // Real y[0..n-1] result
4   {
5       nh := n/2
6       y[0]  := x[0]
7       y[nh] := x[nh]
8       phi := PI/2/n
9       for k:=1 to nh-1
10      {
11          c := cos(phi*k)
12          s := sin(phi*k)
13
14          cps := (c+s)*sqrt(1/2)
15          cms := (c-s)*sqrt(1/2)
16
17          y[k]   := cms*x[k] + cps*x[n-k]
18          y[n-k] := cps*x[k] - cms*x[n-k]
19      }
20  }
```

The C++ equivalent is [FXT: `cos_rot()` in dctdst/cosrot.cc].

Pseudo code for the computation of the DCT via FHT:

```
1   procedure dcth(x[], ldn)
2   // real x[0..n-1] input,result
3   {
4       n := 2**ldn
5       real y[0..n-1]  // workspace
6       unzip_rev(x, y, n)
7       fht(y[],ldn)
8       cos_rot(y[], x[], n)
9   }
```

where `unzip_rev()` is the reversed unzip permutation (see section 2.6 on page 103):

```
1   procedure unzip_rev(a[], b[], n)
2   // real a[0..n-1] input
3   // real b[0..n-1] result
4   {
5       nh := n/2
6       for k:=0 to nh-1
7       {
8           k2 := 2*k
9           b[k]    := a[k2]
10          b[nh+k] := a[n-1-k2]
11      }
12  }
```

Pseudo code for the computation of the inverse discrete cosine transform via FHT:

```
procedure idcth(x[], ldn)
// real x[0..n-1] input,result
{
    n := 2**ldn
    real y[0..n-1]  // workspace
    cos_rot(x[], y[], n);
    fht(y[],ldn)
    zip_rev(y[], x[], n)
}
```

where the routine `zip_rev()` is the reversed zip permutation:

```
procedure zip_rev(a[], b[], n)
// real a[0..n-1] input
// real b[0..n-1] result
{
    nh := n/2
    for k:=0 to nh-1
    {
        k2 := 2*k
        b[k]    := a[k2]
        b[nh+k] := a[n-1-k2]
    }
}
```

The C++ implementations of both the forward and the backward transform [FXT: dctdst/dcth.cc] avoid the temporary array if no scratch space is supplied. The algorithms are given in [194] and [195]. An alternative variant for the computation of the DCT that also uses the FHT is given in [FXT: dctdst/dctzapata.cc], the algorithm is described in [16].

24.12 Discrete sine transform (DST) by DCT

The basis of the *discrete sine transform* (DST) is

$$u(k) \quad = \quad \sin\left(\frac{\pi\,(k+1)\,(i+1/2)}{n}\right) \tag{24.12-1}$$

Pseudo code for the computation of the DST via the discrete cosine transform (DCT):

```
procedure dst(x[],ldn)
// real x[0..n-1] input,result
{
    n := 2**ldn
    nh := n/2

    for k:=1 to n-1 step 2
    {
        x[k] := -x[k]
    }

    dct(x, ldn)

    for k:=0 to nh-1
    {
        swap(x[k], x[n-1-k])
    }
}
```

The corresponding C++ implementation is [FXT: `dsth()` in dctdst/dsth.cc]. Pseudo code for the computation of the inverse sine transform using the inverse cosine transform:

```
procedure idst(x[],ldn)
// real x[0..n-1] input,result
{
    n := 2**ldn
    nh := n/2

    for k:=0 to nh-1
    {
        swap(x[k], x[n-1-k])
    }
```

```
11
12      idct(x, ldn)
13
14      for k:=1 to n-1 step 2
15      {
16          x[k] := -x[k]
17      }
18  }
```

The C++ version is [FXT: `idsth()` in dctdst/dsth.cc].

24.13 Automatic generation of transform code

FFT *generators* are programs that output FFT routines, usually for fixed (short) lengths. In fact the thoughts here are not at all restricted to FFT codes. However, fast transforms and routines that can be unrolled like those for matrix multiplication or convolution are prime candidates for automated generation.

One can write code generators that have a built-in algorithmic knowledge. We restrict our attention to a simpler method known as *partial evaluation.* Writing such a program is easy: take an existing FFT and change all computations into print statements that emit the necessary code. The process, however, is less than delightful and error-prone.

It would be much better to have another program that takes the existing FFT code as input and emit the code for the generator. Let us call this a *meta-generator*. Implementing such a meta-generator of course is highly nontrivial. It actually is equivalent to writing an interpreter for the language used plus the necessary data flow analysis. A practical compromise is to write a program that, while theoretically not even close to a meta-generator, creates output that, after a little hand editing, is a usable generator code.

One may further want to print the current values of the loop variables of the original code as comments at the beginning of a block. Thereby it is possible to locate the corresponding part (with respect to both file and temporal location) of a piece of generated code in the original file. In addition one may keep the comments of the original code.

With FFTs it is necessary to identify ('reverse engineer') the trigonometric values that occur in the process in terms of the corresponding argument (rational multiples of π). The actual values should be inlined to some greater precision than actually needed, thereby one avoids the generation of multiple copies of the (logically) same value with differences only due to numeric inaccuracies. Printing the arguments, both as they appear and in lowest terms, inside comments helps to understand (or further optimize) the generated code:

```
double c1=.98078528040323044912618223613 4; // == cos(Pi*1/16) == cos(Pi*1/16)
double s1=.19509032201612826784828486847 6; // == sin(Pi*1/16) == sin(Pi*1/16)
double c2=.92387953251128675612818318939 7; // == cos(Pi*2/16) == cos(Pi*1/8)
double s2=.38268343236508977172845998402 9; // == sin(Pi*2/16) == sin(Pi*1/8)
```

Automatic verification of the generated codes against the original is a mandatory part of the process.

A level of abstraction for the array indices is of great use: when the print statements in the generator emit some function of the index instead of its plain value it is easy to generate modified versions of the code for permuted input. That is, instead of

```
cout << "sumdiff(f0, f2, g[" << k0 << "], g[" << k2 << "]);"  << endl;
cout << "sumdiff(f1, f3, g[" << k1 << "], g[" << k3 << "]);"  << endl;
```

use

```
cout << "sumdiff(f0, f2, " << idxf(g,k0) << ", " << idxf(g,k2) << ");"  << endl;
cout << "sumdiff(f1, f3, " << idxf(g,k1) << ", " << idxf(g,k3) << ");"  << endl;
```

where `idxf(g, k)` can be defined to print a modified (for example, revbin-permuted) index `k`.

A generated length-8 DIT FHT core (from [FXT: fht/shortfhtditcore.h]) shall serve as an example:

```
template <typename Type>
inline void fht_dit_core_8(Type *f)
// unrolled version for length 8
{
 { // start initial loop
  { // fi = 0  gi = 1
   Type g0, f0, f1, g1;
   sumdiff(f[0], f[1], f0, g0);
   sumdiff(f[2], f[3], f1, g1);
   sumdiff(f0, f1);
   sumdiff(g0, g1);
   Type s1, c1, s2, c2;
   sumdiff(f[4], f[5], s1, c1);
   sumdiff(f[6], f[7], s2, c2);
   sumdiff(s1, s2);
   sumdiff(f0, s1, f[0], f[4]);
   sumdiff(f1, s2, f[2], f[6]);
   c1 *= SQRT2;
   c2 *= SQRT2;
   sumdiff(g0, c1, f[1], f[5]);
   sumdiff(g1, c2, f[3], f[7]);
  }
 } // end initial loop
}
// opcount by generator:  #mult=2=0.25/pt   #add=22=2.75/pt
```

Generated DIF FHT codes for lengths up to 64 are given in [FXT: fht/shortfhtdifcore.h].

The generated codes can be useful when one wants to spot parts of the original code that need further optimization. Especially repeated trigonometric values and unused symmetries tend to be apparent in the unrolled code.

It is a good idea to let the generator count the number of operations (multiplications, additions, loads and stores) of the code it emits. It is even better if those numbers are compared to the corresponding values found in the compiled assembler code.

Checking the generated machine code

It is possible to have GCC produce the assembler code with the original source interlaced. This is a great tool for code optimization. The necessary commands are (include- and warning flags omitted)

```
# create assembler code:
c++ -S -fverbose-asm -g -O2  test.cc -o test.s
# create asm interlaced with source lines:
as -alhnd test.s  > test.lst
```

For example, the generated length-4 DIT FHT core from [FXT: fht/shortfhtditcore.h] is

```
template <typename Type>
inline void fht_dit_core_4(Type *f)
// unrolled version for length 4
{
   Type f0, f1, f2, f3;
   sumdiff(f[0], f[1], f0, f1);
   sumdiff(f[2], f[3], f2, f3);
   sumdiff(f0, f2, f[0], f[2]);
   sumdiff(f1, f3, f[1], f[3]);
}
```

With `Type` set to `double` the generated assembler is (some editing for readability)

```
     void fht_dit_core_4(double *f)
     {
         double f0, f1, f2, f3;
         sumdiff(f[0], f[1], f0, f1);
               movlpd  (%rdi), %xmm1   #* f, tmp63
               movlpd  8(%rdi), %xmm0  #, tmp64
         sumdiff(f[2], f[3], f2, f3);
               movlpd  16(%rdi), %xmm2 #, tmp67
               movsd   %xmm1, %xmm3    # tmp63, f0
               subsd   %xmm0, %xmm1    # tmp64, f1
```

```
11                movsd   %xmm2, %xmm4    # tmp67, f2
12                addsd   %xmm0, %xmm3    # tmp64, f0
13                movlpd  24(%rdi), %xmm0 #, tmp68
14                addsd   %xmm0, %xmm4    # tmp68, f2
15                subsd   %xmm0, %xmm2    # tmp68, f3
16          sumdiff(f0, f2, f[0], f[2]);
17                movsd   %xmm3, %xmm0    # f0, tmp71
18                addsd   %xmm4, %xmm0    # f2, tmp71
19                subsd   %xmm4, %xmm3    # f2, f0
20                movsd   %xmm0, (%rdi)   # tmp71,* f
21          sumdiff(f1, f3, f[1], f[3]);
22                movsd   %xmm1, %xmm0    # f1, tmp73
23                subsd   %xmm2, %xmm1    # f3, f1
24                movsd   %xmm3, 16(%rdi) # f0,
25                addsd   %xmm2, %xmm0    # f3, tmp73
26                movsd   %xmm1, 24(%rdi) # f1,
27                movsd   %xmm0, 8(%rdi)  # tmp73,
28      }
```

Note that the assembler code is not always in sync with the corresponding source lines, especially with higher levels of optimization.

24.14 Eigenvectors of the Fourier and Hartley transform *

Let $a_S := a + \overline{a}$ be the symmetric part of a sequence a, then

$$\mathcal{F}\big[\,\mathcal{F}\big[\,a_S\,\big]\,\big] \;=\; a_S \tag{24.14-1}$$

Now let $u_+ := a_S + \mathcal{F}\big[\,a_S\,\big]$ and $u_- := a_S - \mathcal{F}\big[\,a_S\,\big]$, then

$$\mathcal{F}\big[\,u_+\,\big] \;=\; \mathcal{F}\big[\,a_S\,\big] + a_S \;=\; a_S + \mathcal{F}\big[\,a_S\,\big] \;=\; +1 \cdot u_+ \tag{24.14-2a}$$
$$\mathcal{F}\big[\,u_-\,\big] \;=\; \mathcal{F}\big[\,a_S\,\big] - a_S \;=\; -(a_S - \mathcal{F}\big[\,a_S\,\big]) \;=\; -1 \cdot u_- \tag{24.14-2b}$$

Both u_+ and u_- are symmetric. For $a_A := a - \overline{a}$, the antisymmetric part of a, we have

$$\mathcal{F}\big[\,\mathcal{F}\big[\,a_A\,\big]\,\big] \;=\; -a_A \tag{24.14-3}$$

Therefore with $v_+ := a_A + i\,\mathcal{F}\big[\,a_A\,\big]$ and $v_- := a_A - i\,\mathcal{F}\big[\,a_A\,\big]$:

$$\mathcal{F}\big[\,v_+\,\big] \;=\; \mathcal{F}\big[\,a_A\,\big] - i\,a_A \;=\; -i\,(a_A + i\,\mathcal{F}\big[\,a_A\,\big]) \;=\; -i \cdot v_+ \tag{24.14-4a}$$
$$\mathcal{F}\big[\,v_-\,\big] \;=\; \mathcal{F}\big[\,a_A\,\big] + i\,a_A \;=\; +i\,(a_A - i\,\mathcal{F}\big[\,a_A\,\big]) \;=\; +i \cdot v_- \tag{24.14-4b}$$

Both v_+ and v_- are antisymmetric. The sequences u_+, u_-, v_+, and v_- are *eigenvectors* of the FT, with *eigenvalues* $+1$, -1, $-i$ and $+i$ respectively. The eigenvectors are pairwise perpendicular. Using the relation

$$a \;=\; \frac{1}{2}\,(u_+ + u_- + v_+ + v_-) \tag{24.14-5}$$

we can, for a given sequence, find a transform that is a 'square root' of the FT: compute u_+, u_-, v_+, and v_-, and a transform $\mathcal{F}^\lambda\,[a]$ for $\lambda \in \mathbb{R}$ as

$$\mathcal{F}^\lambda\,[a] \;=\; \frac{1}{2}\Big((+1)^\lambda\,u_+ + (-1)^\lambda\,u_- + (-i)^\lambda\,v_+ + (+i)^\lambda\,v_-\Big) \tag{24.14-6}$$

This transform is called the *fractional (order) Fourier transform* (or simply fractional FT, but see section 21.5.3 on page 442). Then $\mathcal{F}^0\,[a]$ is the identity and $\mathcal{F}^1\,[a]$ is the usual FT. The transform $\mathcal{F}^{1/2}\,[a]$ is a transform so that $\mathcal{F}^{1/2}\,\big[\mathcal{F}^{1/2}\,[a]\big] = \mathcal{F}\big[\,a\,\big]$, that is, a 'square root' of the FT. The transform $\mathcal{F}^{1/2}\,[a]$ is not unique as the expressions $\pm 1^{1/2}$ and $\pm i^{1/2}$ are not.

The eigenvectors of the Hartley Transform are

$$u_+ := a + \mathcal{H}[a] \tag{24.14-7a}$$
$$u_- := a - \mathcal{H}[a] \tag{24.14-7b}$$

The eigenvalues are ± 1, one has $\mathcal{H}[u_+] = +1 \cdot u_+$ and $\mathcal{H}[u_-] = -1 \cdot u_-$.

Let M be the $n \times n$ matrix corresponding to the length-n Fourier transform with positive sign σ, that is $M_{r,c} = 1/\sqrt{n} \exp(2\pi i r c/n)$. Then its characteristic polynomial (see relation 40.5-2 on page 889) is

$$p(x) = (x-1)^{\lfloor (n+4)/4 \rfloor} (x+1)^{\lfloor (n+2)/4 \rfloor} (x-i)^{\lfloor (n+1)/4 \rfloor} (x+i)^{\lfloor (n-1)/4 \rfloor} \tag{24.14-8}$$

We write $p(x) = x^n + c_{n-1} x^{n-1} + \ldots + c_1 x + c_0$. The trace of the matrix M is

$$\mathrm{Tr}(M) = \sqrt{n} \sum_{k=0}^{n-1} \exp\left(2 i \pi k^2/n\right) \tag{24.14-9}$$

It equals ($-c_{n-1}$, the negative sum of all roots of $p(x)$, and)

$$1+i,\ +1,\ 0,\ +i \tag{24.14-10}$$

for $n \bmod 4 \equiv 0, 1, 2, 3$, respectively. A closed form is $(1 + i^{-n}) / (1 - i)$. The generating function for the sequence of values is $((1+i) - x) / \left(1 + (-1+i)\, x - i\, x^2\right)$.

The determinant of M equals ($(-1)^n c_0$, $(-1)^n$ times the product of all roots of $p(x)$, and)

$$+i,\ +1,\ -1,\ -i,\ -i,\ -1,\ +1,\ +i \tag{24.14-11}$$

for $n \bmod 8 \equiv 0, 1, 2, \ldots, 7$. A closed form is $(1+i)/2\ (1 + (-i)^n)$. The generating function for the sequence is $\left(i + x - x^2 - i\, x^3\right) / \left(1 + x^4\right)$.

Chapter 25

Number theoretic transforms (NTTs)

We introduce the number theoretic transforms (NTTs). After understanding the necessary concepts from number theory (see also chapter 37) it turns out that the routines for the fast NTTs are rather straightforward translations of the FFT algorithms. We give radix-2 and radix-4 routines but there should be no difficulty to translate any given (complex valued) FFT algorithm into the equivalent NTT algorithm. For the translation of real valued FFT (or FHT) routines one needs to express sines and cosines in modular arithmetic, this is presented in sections 37.12.6 and 37.12.7.

As no rounding errors occur with the underlying modular arithmetic the main application of NTTs is the fast computation of exact convolutions.

25.1 Prime moduli for NTTs

How to make a number theoretic transform out of your FFT:
'*Replace* $exp(\pm 2\pi i/n)$ *by a primitive n-th root of unity, done.*'

We want to implement FFTs in $\mathbb{Z}/m\mathbb{Z}$ (the ring of integers modulo some integer m) instead of $\mathbb{C}$, the (field of the) complex numbers. These FFTs are called *number theoretic transforms* (NTTs), *mod m FFTs* or (if m is a prime) *prime modulus transforms*.

There is a restriction for the choice of m: for a length n NTT we need a primitive n-th root of unity. A number r is called an n-th root of unity if $r^n = 1$. It is called a *primitive n-th root* if $r^k \neq 1 \; \forall \, k < n$.

In $\mathbb{C}$ matters are simple: $e^{\pm 2\pi i/n}$ is a primitive n-th root of unity for arbitrary n. For example, $e^{2\pi i/21}$ is a primitive 21-th root of unity. Now $r = e^{2\pi i/3}$ is also 21-th root of unity but not a primitive root, because $r^3 = 1$. A primitive n-th root of 1 in $\mathbb{Z}/m\mathbb{Z}$ is also called an *element of order n*. The 'cyclic' property of the elements r of order n lies in the heart of all FFT algorithms: $r^{n+k} = r^k$.

In $\mathbb{Z}/m\mathbb{Z}$ things are not that simple: for a given modulus m primitive n-th roots of unity do not exist for arbitrary n. They exist for some maximal order R only. Roots of unity of an order different from R are available only for the divisors d_i of R: r^{R/d_i} is a d_i-th root of unity because $(r^{R/d_i})^{d_i} = r^R = 1$.

Therefore n, the length of the transform, must divide the maximal order R. This is the first condition for NTTs.

The operations needed in FFTs are addition, subtraction and multiplication. Division is not needed, except for division by n for the final normalization after transform and back-transform. Division by n is multiplication by the inverse of n, so n must be invertible in $\mathbb{Z}/m\mathbb{Z}$.

Therefore n, the length of the transform, must be coprime to the modulus m: $\gcd(n, m) = 1$. This is the second condition for NTTs.

We restrict our attention to prime moduli, though NTTs are also possible with composite moduli. If the modulus is a prime p then $\mathbb{Z}/p\mathbb{Z}$ is the field $\mathbb{F}_p = \mathrm{GF}(p)$: all elements except 0 have inverses and 'division is possible'. Thereby the second condition is trivially fulfilled for all NTT lengths $n < p$: a prime p is coprime to all integers $n < p$.

Roots of unity are available for the maximal order $R = p-1$ and its divisors: Therefore the first condition on n for a length-n mod p NTT being possible is that n divides $p - 1$. This restricts the choice for p to primes of the form $p = v\,n + 1$: for length-$n = 2^k$ NTTs one will use primes like $p = 3 \cdot 5 \cdot 2^{27} + 1$ (31 bits), $p = 13 \cdot 2^{28} + 1$ (32 bits), $p = 3 \cdot 29 \cdot 2^{56} + 1$ (63 bits) or $p = 27 \cdot 2^{59} + 1$ (64 bits).

```
arg 1: 62 == wb  [word bits, wb<=63]  default=62
arg 2: 0.01 == deltab  [results are in the range [wb-deltab, wb]]  default=0.01
 minb = 61.99 = wb-0.01
arg 3: 44 == minx  [log_2(min(fftlen))]  default=44
  ---- x = 44: -----
4580495072570638337 = 0x3f91300000000001 = 1 + 2^44 * 83 * 3137          (61.9902 bits)
4581058022524059649 = 0x3f93300000000001 = 1 + 2^44 * 3 * 11 * 13 * 607  (61.9904 bits)
4582113553686724609 = 0x3f96f00000000001 = 1 + 2^44 * 3 * 7 * 79 * 157   (61.9907 bits)
4585702359639785473 = 0x3fa3b00000000001 = 1 + 2^44 * 3^2 * 11 * 2633    (61.9918 bits)
4587039365779161089 = 0x3fa8700000000001 = 1 + 2^44 * 7 * 193^2          (61.9923 bits)
4587391209500049409 = 0x3fa9b00000000001 = 1 + 2^44 * 3 * 17 * 5113      (61.9924 bits)
4588130081313914881 = 0x3fac500000000001 = 1 + 2^44 * 3 * 5 * 17387      (61.9926 bits)
4589572640569556993 = 0x3fb1700000000001 = 1 + 2^44 * 11 * 37 * 641      (61.9931 bits)
[--snip--]
4610999923171655681 = 0x3ffd900000000001 = 1 + 2^44 * 5 * 19 * 31 * 89   (61.9998 bits)
4611105476287922177 = 0x3ffdf00000000001 = 1 + 2^44 * 262111             (61.9998 bits)
  ---- x = 45: -----
4580336742896238593 = 0x3f90a00000000001 = 1 + 2^45 * 29 * 67^2          (61.9902 bits)
4581533011547258881 = 0x3f94e00000000001 = 1 + 2^45 * 3 * 5 * 8681       (61.9905 bits)
4584347761314365441 = 0x3f9ee00000000001 = 1 + 2^45 * 5 * 11 * 23 * 103  (61.9914 bits)
4587655092290715649 = 0x3faaa00000000001 = 1 + 2^45 * 3 * 7^2 * 887      (61.9925 bits)
[--snip--]
  ---- x = 48: -----
4585508845593296897 = 0x3fa3000000000001 = 1 + 2^48 * 11 * 1481          (61.9918 bits)
  ---- x = 49: -----
4582975570802900993 = 0x3f9a000000000001 = 1 + 2^49 * 7 * 1163           (61.991 bits)
4595360469778169857 = 0x3fc6000000000001 = 1 + 2^49 * 3^2 * 907          (61.9949 bits)
  ---- x = 50: -----
4601552919265804289 = 0x3fdc000000000001 = 1 + 2^50 * 61 * 67            (61.9968 bits)
```

Figure 25.1-A: Primes suitable for NTTs of lengths dividing 2^{44}.

```
     modulus (hex)       ==    factorization + 1     log(m-1)/log(2)
    0x3f40f80000000001  == 2^43.3^2.5^2.7^2.47+1         61.9831
    0x3c0eb50000000001  == 2^40.3^3.5^2.7^3.17+1         61.9083
    0x3d673d0000000001  == 2^40.3^2.5^3.7^2.73+1         61.9402
    0x3fc22b0000000001  == 2^40.3^2.5^2.7^2.379+1        61.9945
    0x3bf6190000000001  == 2^40.3^2.5^3.7.499+1          61.906
    0x3d1d690000000001  == 2^40.3^2.5^2.7.2543+1         61.9335
    0x3d8c270000000001  == 2^40.3^2.5^2.7.13.197+1       61.9436
    0x3e8e8d0000000001  == 2^40.3^2.5^2.7.19.137+1       61.9671
    0x3ee4af0000000001  == 2^40.3^2.5^2.7.2617+1         61.9748
    0x3ed23a0000000001  == 2^41.3^2.5^2.7.1307+1         61.9732
    0x3fafb60000000001  == 2^41.3^2.5^4.7.53+1           61.9929
    0x3c46140000000001  == 2^42.3^3.5^2.7.11.19+1        61.9135
    0x3e32440000000001  == 2^42.3^2.5^2.7.647+1          61.9588
    0x3d23900000000001  == 2^44.3^3.5^2.7.53+1           61.934
```

Figure 25.1-B: Primes suitable for NTTs of lengths dividing $2^{40}\,3^2\,5^2\,7$.

Primes suitable with NTTs (sometimes called *FFT-primes*) can be generated with the program [FXT: mod/fftprimes-demo.cc]. A shortened sample output is shown in figure 25.1-A A few moduli that allow for transforms of lengths dividing $2^{40} \cdot 3^2 \cdot 5^2 \cdot 7$ are shown in figure 25.1-B, the data is taken from [FXT: mod/moduli.txt]. We note that primality of moduli suitable for NTTs can easily by tested using Proth's theorem, see section 37.11.3.1 on page 784.

25.2 Implementation of NTTs

To implement NTTs (modulo m, length n) one has to implement modular arithmetics and replace $e^{\pm 2\pi i/n}$ by an primitive n-th root r of unity in $\mathbb{Z}/m\mathbb{Z}$ in the code. A C++ class implementing modular arithmetics in FXT is [FXT: `class mod` in mod/mod.h].

For the inverse transform one uses the (mod m) inverse r^{-1} of r that was used for the forward transform. The element r^{-1} is also a primitive n-th root. Methods for the computation of the modular inverse are described in section 37.1.4 on page 756 (*gcd* algorithm) and on page 768 (powering algorithm: $r^{-1} = r^{R-1}$).

While the notion of the Fourier transform as a 'decomposition into frequencies' seems to be meaningless for NTTs the algorithms are denoted with 'decimation in time/frequency' in analogy to those in the complex domain.

The nice feature of NTTs is that there is no loss of precision in the transform as with the floating point FFTs. Using the trigonometric recursion in its most naive form is mandatory, as the computation of roots of unity is expensive.

25.2.1 Radix-2 DIT NTT

Pseudo code for the radix-2 decimation in time (DIT) NTT (to be called with `ldn=log2(n)`):

```
1   procedure mod_fft_dit2(f[], ldn, is)
2   // mod_type f[0..2**ldn-1]
3   {
4       n := 2**ldn
5
6       rn := element_of_order(n)  // (mod_type)
7
8       if is<0 then  rn := rn**(-1)
9
10      revbin_permute(f[], n)
11
12      for ldm:=1 to ldn
13      {
14          m  := 2**ldm
15          mh := m/2
16
17          dw := rn**(2**(ldn-ldm))   // (mod_type)
18          w  := 1                    // (mod_type)
19
20          for j:=0 to mh-1
21          {
22              for r:=0 to n-m step m
23              {
24                  t1 := r+j
25                  t2 := t1+mh
26
27                  v := f[t2]*w  // (mod_type)
28                  u := f[t1]    // (mod_type)
29
30                  f[t1] := u+v
31                  f[t2] := u-v
32              }
33
34              w := w*dw  // trig recursion
35          }
36      }
37  }
```

As shown in section 20.3.1 on page 396 it is a good idea to extract the `ldm==1` stage of the outermost loop: Replace

```
for ldm:=1 to ldn
{
```

by

```
for r:=0 to n-1 step 2
{
    {f[r], f[r+1]} := {f[r]+f[r+1], f[r]-f[r+1]}
```

```
    }
    for ldm:=2 to ldn
    {
```

The C++ implementation is given in [FXT: ntt/nttdit2.cc]:

```
void
ntt_dit2_core(mod *f, ulong ldn, int is)
// Auxiliary routine for ntt_dit2()
// Decimation in time (DIT) radix-2 FFT
// Input data must be in revbin_permuted order
// ldn := base-2 logarithm of the array length
// is := sign of the transform
{
    const ulong n = 1UL<<ldn;

    for (ulong i=0; i<n; i+=2)  sumdiff(f[i], f[i+1]);

    for (ulong ldm=2; ldm<=ldn; ++ldm)
    {
        const ulong m = (1UL<<ldm);
        const ulong mh = (m>>1);

        const mod dw = mod::root2pow( is>0 ? ldm : -ldm );
        mod w = (mod::one);

        for (ulong j=0; j<mh; ++j)
        {
            for (ulong r=0; r<n; r+=m)
            {
                const ulong t1 = r + j;
                const ulong t2 = t1 + mh;

                mod v = f[t2] * w;
                mod u = f[t1];

                f[t1] = u + v;
                f[t2] = u - v;
            }
            w *= dw;
        }
    }
}
```

```
void
ntt_dit2(mod *f, ulong ldn, int is)
// Radix-2 decimation in time (DIT) NTT
{
    revbin_permute(f, 1UL<<ldn);
    ntt_dit2_core(f, ldn, is);
}
```

The elements of order 2^k are precomputed upon initialization of the `mod` class. The call to `mod::root2pow()` is a simple table lookup.

25.2.2 Radix-2 DIF NTT

Pseudo code for the radix-2 decimation in frequency (DIF) NTT:

```
procedure mod_fft_dif2(f[], ldn, is)
// mod_type f[0..2**ldn-1]
{
    n := 2**ldn
    dw := element_of_order(n)  // (mod_type)

    if is<0 then  dw := rn**(-1)

    for  ldm:=ldn to 1 step -1
    {
        m  := 2**ldm
        mh := m/2

        w := 1  // (mod_type)

        for j:=0 to mh-1
```

```
17          {
18              for r:=0 to n-m step m
19              {
20                  t1 := r+j
21                  t2 := t1+mh
22
23                  v := f[t2]  // (mod_type)
24                  u := f[t1]  // (mod_type)
25
26                  f[t1] :=  u+v
27                  f[t2] := (u-v)*w
28              }
29
30              w := w*dw  // trig recursion
31          }
32
33          dw := dw*dw
34      }
35
36      revbin_permute(f[], n)
37  }
```

As in section 20.3.2 on page 399 extract the `ldm==1` stage of the outermost loop: Replace the line

```
for  ldm:=ldn to 1 step -1
```

by

```
for  ldm:=ldn to 2 step -1
```

and insert

```
for r:=0 to n-1 step 2
{
    {f[r], f[r+1]} := {f[r]+f[r+1], f[r]-f[r+1]}
}
```

before the call of `revbin_permute(f[],n)`.

The C++ implementation is given in [FXT: ntt/nttdif2.cc]:

```
1   void
2   ntt_dif2_core(mod *f, ulong ldn, int is)
3   // Auxiliary routine for ntt_dif2().
4   // Decimation in frequency (DIF) radix-2 NTT.
5   // Output data is in revbin_permuted order.
6   // ldn := base-2 logarithm of the array length.
7   // is := sign of the transform
8   {
9       const ulong n = (1UL<<ldn);
10      mod dw = mod::root2pow( is>0 ? ldn : -ldn );
11
12      for (ulong ldm=ldn; ldm>1; --ldm)
13      {
14          const ulong m = (1UL<<ldm);
15          const ulong mh = (m>>1);
16
17          mod w = mod::one;
18
19          for (ulong j=0; j<mh; ++j)
20          {
21              for (ulong r=0; r<n; r+=m)
22              {
23                  const ulong t1 = r + j;
24                  const ulong t2 = t1 + mh;
25
26                  mod v = f[t2];
27                  mod u = f[t1];
28
29                  f[t1] = (u + v);
30                  f[t2] = (u - v) * w;
31              }
32              w *= dw;
33          }
34          dw *= dw;
35      }
36
37      for (ulong i=0; i<n; i+=2)  sumdiff(f[i], f[i+1]);
```

```
}

void
ntt_dif2(mod *f, ulong ldn, int is)
// Radix-2 decimation in frequency (DIF) NTT
{
    ntt_dif2_core(f, ldn, is);
    revbin_permute(f, 1UL<<ldn);
}
```

25.2.3 Radix-4 NTTs

The radix-4 versions of the NTT are straightforward translations of the routines that use complex numbers. We simply give the C++ implementations

25.2.3.1 Decimation in time (DIT) algorithm

Code for a radix-4 decimation in time (DIT) NTT [FXT: ntt/nttdit4.cc]:

```
static const ulong LX = 2;

void
ntt_dit4_core(mod *f, ulong ldn, int is)
// Auxiliary routine for ntt_dit4()
// Decimation in time (DIT) radix-4 NTT
// Input data must be in revbin_permuted order
// ldn := base-2 logarithm of the array length
// is := sign of the transform
{
    const ulong n = (1UL<<ldn);

    if ( ldn & 1 )  // n is not a power of 4, need a radix-2 step
    {
        for (ulong i=0; i<n; i+=2)  sumdiff(f[i], f[i+1]);
    }

    const mod imag = mod::root2pow( is>0 ? 2 : -2 );

    ulong ldm = LX + (ldn&1);
    for ( ; ldm<=ldn ; ldm+=LX)
    {
        const ulong m = (1UL<<ldm);
        const ulong m4 = (m>>LX);

        const mod dw = mod::root2pow( is>0 ? ldm : -ldm );
        mod w = (mod::one);
        mod w2 = w;
        mod w3 = w;

        for (ulong j=0; j<m4; j++)
        {
            for (ulong r=0, i0=j+r; r<n; r+=m, i0+=m)
            {
                const ulong i1 = i0 + m4;
                const ulong i2 = i1 + m4;
                const ulong i3 = i2 + m4;

                mod a0 = f[i0];
                mod a2 = f[i1] * w2;
                mod a1 = f[i2] * w;
                mod a3 = f[i3] * w3;

                mod t02 = a0 + a2;
                mod t13 = a1 + a3;

                f[i0] = t02 + t13;
                f[i2] = t02 - t13;

                t02 = a0 - a2;
                t13 = a1 - a3;
                t13 *= imag;

```

```
                f[i1] = t02 + t13;
                f[i3] = t02 - t13;
            }

            w *= dw;
            w2 = w * w;
            w3 = w * w2;
        }
    }
}
```

```
void
ntt_dit4(mod *f, ulong ldn, int is)
// Radix-4 decimation in time (DIT) NTT
{
    revbin_permute(f, 1UL<<ldn);
    ntt_dit4_core(f, ldn, is);
}
```

25.2.3.2 Decimation in frequency (DIF) algorithm

Code for a radix-4 decimation in frequency (DIT) NTT [FXT: ntt/nttdif4.cc]:

```
static const ulong LX = 2;

void
ntt_dif4_core(mod *f, ulong ldn, int is)
// Auxiliary routine for ntt_dif4().
// Decimation in frequency (DIF) radix-4 NTT.
// Output data is in revbin_permuted order.
// ldn := base-2 logarithm of the array length.
// is := sign of the transform
{
    const ulong n = (1UL<<ldn);

    const mod imag = mod::root2pow( is>0 ? 2 : -2 );

    for (ulong ldm=ldn; ldm>=LX; ldm-=LX)
    {
        const ulong m = (1UL<<ldm);
        const ulong m4 = (m>>LX);

        const mod dw = mod::root2pow( is>0 ? ldm : -ldm );
        mod w = (mod::one);
        mod w2 = w;
        mod w3 = w;

        for (ulong j=0; j<m4; j++)
        {
            for (ulong r=0, i0=j+r; r<n; r+=m, i0+=m)
            {
                const ulong i1 = i0 + m4;
                const ulong i2 = i1 + m4;
                const ulong i3 = i2 + m4;

                mod a0 = f[i0];
                mod a1 = f[i1];
                mod a2 = f[i2];
                mod a3 = f[i3];

                mod t02 = a0 + a2;
                mod t13 = a1 + a3;

                f[i0] = (t02 + t13);
                f[i1] = (t02 - t13) * w2;

                t02 = a0 - a2;
                t13 = a1 - a3;
                t13 *= imag;

                f[i2] = (t02 + t13) * w;
                f[i3] = (t02 - t13) * w3;
            }

            w *= dw;
            w2 = w * w;
```

```
54              w3 = w * w2;
55          }
56      }
57
58      if ( ldn & 1 )  // n is not a power of 4, need a radix-2 step
59      {
60          for (ulong i=0; i<n; i+=2)  sumdiff(f[i], f[i+1]);
61      }
62  }
```

```
1   void
2   ntt_dif4(mod *f, ulong ldn, int is)
3   // Radix-4 decimation in frequency (DIF) NTT
4   {
5       ntt_dif4_core(f, ldn, is);
6       revbin_permute(f, 1UL<<ldn);
7   }
```

25.3 Convolution with NTTs

The NTTs are natural candidates for the computation of exact integer convolutions, as used in high precision multiplication algorithms. One must keep in mind that all computations are modulo m, the largest value that can be represented is $m-1$. Choosing a modulus that is greater than the maximal possible value of the result avoids any truncation. If m does not fit into a single machine word this may slow down the computation unacceptably.

It is better to choose m as the product of several coprime moduli m_i that are all just below machine word size, compute the convolutions for each modulus m_i, and finally use the Chinese Remainder Theorem (see section 37.7 on page 769) to obtain the result modulo m.

If length-n FFTs are used for convolution there must be an inverse element for n. This imposes the condition $\gcd(n, m) = 1$, the modulus must be coprime to n. For length-2^k FFTs this simply means that m must be odd.

C++ code for the NTT based exact convolution can be found in [FXT: ntt/nttcnvl.cc]. The routines are virtually identical to their complex equivalents. For example, a routine for cyclic self-convolution can be given as

```
1   void
2   ntt_auto_convolution(mod *f, ulong ldn)
3   // Cyclic (self-)convolution.
4   // Use zero padded data for linear convolution.
5   {
6       assert_two_invertible();  // so we can normalize later
7
8       const int is = +1;
9
10      ntt_dif4_core(f, ldn, is);  // transform
11
12      const ulong n = (1UL<<ldn);
13      for (ulong i=0; i<n; ++i)  f[i] *= f[i];  // multiply element-wise
14
15      ntt_dit4_core(f, ldn, -is);  // inverse transform
16
17      multiply_val(f, n, (mod(n)).inv() );  // normalize
18  }
```

The revbin permutations are avoided as explained in section 21.1.3 on page 429.

For further applications of the NTT see the survey article [138] and the references given there.

Chapter 26

Fast wavelet transforms

The discrete wavelet transforms are a class of transforms that can be computed in linear time. We treat wavelet transforms whose basis functions have compact support. These can be derived as a generalization of the Haar transform.

26.1 Wavelet filters

We motivate the *wavelet transform* as a generalization of the 'standard' Haar transform given in section 23.1 on page 483. We reformulate the Haar transform as a sequence of filtering steps.

We consider only (moving average) filters F defined by n coefficients ('taps') f_0, f_1, ..., f_{n-1}. Let A be the length-N sequence a_0, a_1, ..., a_{N-1}. We define $F_k(A)$ as the weighted sum

$$F_k(A) \quad := \quad \sum_{j=0}^{n-1} f_j \, a_{k+j \bmod N} \tag{26.1-1}$$

That is, $F_k(A)$ is the result of applying the filter F to the n elements a_k, a_{k+1}, a_{k+2}, ... a_{k+n-1}, possibly wrapping around.

Now assume that N is a power of two. Let H be the low-pass filter defined by $h_0 = h_1 = +1/\sqrt{2}$, and G be the high-pass filter defined by $g_0 = +1/\sqrt{2}$, $g_1 = -1/\sqrt{2}$. A single filtering step of the Haar transform consists of

- Computing the sums $s_0 = H_0(A)$, $s_2 = H_2(A)$, $s_4 = H_4(4)$, ..., $s_{N-2} = H_{N-2}(A)$
- Computing the differences $d_0 = G_0(A)$, $d_2 = G_2(A)$, $d_4 = G_4(4)$, ..., $d_{N-2} = G_{N-2}(A)$
- Writing the sums to the left half of A, and the differences to the right half:
  $A = [s_0, s_2, s_4, s_6, \ldots, s_{N-2}, d_0, d_2, d_4, d_6, \ldots, d_{N-2}]$

The Haar transform is obtained by applying the step to the whole sequence, then to its left half, then to its left quarter, ..., the left four elements, the left two elements. With the Haar transform no wrap-around occurs.

A the analogous filtering step for the wavelet transform is obtained by defining two length-n filters H (low-pass) and G (high-pass) subject to certain conditions. Firstly, we consider only filters with an even number n of coefficients.

Secondly, we define coefficients of G to be the reversed sequence of the coefficients of H with alternating signs:
$g_0 = +h_{n-1}$, $g_1 = -h_{n-2}$, $g_2 = +h_{n-3}$, $g_4 = -h_{n-3}$, ..., $g_{n-3} = -h_2$, $g_{n-2} = +h_1$, $g_{n-1} = -h_0$.

Thirdly, we require that the resulting transform is orthogonal. Let S be the matrix corresponding to one filtering step, ignoring the order:

$$S\,A \quad = \quad [s_0, d_0, s_2, d_2, s_4, d_4, s_6, d_6, \ldots, s_{N-2}, d_{N-2}] \tag{26.1-2}$$

With length-6 filters and $N = 16$ the matrix S would be

$$S \quad = \quad \begin{bmatrix}
h_0 & h_1 & h_2 & h_3 & h_4 & h_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
g_0 & g_1 & g_2 & g_3 & g_4 & g_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & h_0 & h_1 & h_2 & h_3 & h_4 & h_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & g_0 & g_1 & g_2 & g_3 & g_4 & g_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & h_0 & h_1 & h_2 & h_3 & h_4 & h_5 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & g_0 & g_1 & g_2 & g_3 & g_4 & g_5 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & h_0 & h_1 & h_2 & h_3 & h_4 & h_5 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & g_0 & g_1 & g_2 & g_3 & g_4 & g_5 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & h_0 & h_1 & h_2 & h_3 & h_4 & h_5 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & g_0 & g_1 & g_2 & g_3 & g_4 & g_5 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & h_0 & h_1 & h_2 & h_3 & h_4 & h_5 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & g_0 & g_1 & g_2 & g_3 & g_4 & g_5 \\
h_4 & h_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & h_0 & h_1 & h_2 & h_3 \\
g_4 & g_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & g_0 & g_1 & g_2 & g_3 \\
h_2 & h_3 & h_4 & h_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & h_0 & h_1 \\
g_2 & g_3 & g_4 & g_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & g_0 & g_1
\end{bmatrix} \tag{26.1-3a}$$

$$= \quad \begin{bmatrix}
+h_0 & +h_1 & +h_2 & +h_3 & +h_4 & +h_5 & 0 & 0 & 0 & 0 & 0 & \ldots & 0 \\
+h_5 & -h_4 & +h_3 & -h_2 & +h_1 & -h_0 & 0 & 0 & 0 & 0 & 0 & \ldots & 0 \\
0 & 0 & +h_0 & +h_1 & +h_2 & +h_3 & +h_4 & +h_5 & 0 & 0 & 0 & \ldots & 0 \\
0 & 0 & +h_5 & -h_4 & +h_3 & -h_2 & +h_1 & -h_0 & 0 & 0 & 0 & \ldots & 0 \\
0 & 0 & 0 & 0 & +h_0 & +h_1 & +h_2 & +h_3 & +h_4 & +h_5 & 0 & \ldots & 0 \\
0 & 0 & 0 & 0 & +h_5 & -h_4 & +h_3 & -h_2 & +h_1 & -h_0 & 0 & \ldots & 0 \\
 & & & \ddots & & & & & & \ddots & & &
\end{bmatrix} \tag{26.1-3b}$$

The orthogonality requires that $S\,S^T = \mathrm{id}$, that is (setting $h_j = 0$ for $j < 0$ and $j \geq n$)

$$\sum_j h_j^2 \quad = \quad 1 \tag{26.1-4a}$$

$$\sum_j h_j\, h_{j+2} \quad = \quad 0 \tag{26.1-4b}$$

$$\sum_j h_j\, h_{j+4} \quad = \quad 0 \tag{26.1-4c}$$

In general, the following $n/2$ *wavelet conditions* are obtained:

$$\sum_j h_j^2 \quad = \quad 1 \tag{26.1-5a}$$

$$\sum_j h_j\, h_{j+2i} \quad = \quad 0 \quad \text{where} \quad i = 1, 2, 3, \ldots, n/2 - 1 \tag{26.1-5b}$$

We call a filter H satisfying these conditions a *wavelet filter*.

For the wavelet transform with $n = 2$ filter taps there is only condition, $h_0^2 + h_1^2 = 1$, leading to the parametric solution $h_0 = \sin(\phi)$, $h_1 = \cos(\phi)$. Setting $\phi = \pi/4$ one obtains $h_0 = h_1 = 1/\sqrt{2}$, corresponding to the Haar transform.

26.2 Implementation

A container class for wavelet filters is [FXT: `class wavelet_filter` in wavelet/waveletfilter.h]:

```
class wavelet_filter
{
public:
    double *h_;  // low-pass filter
    double *g_;  // high-pass filter
    ulong n_;    // number of taps

    void ctor_core()
    {
        h_ = new double[n_];
        g_ = new double[n_];
    }

    wavelet_filter(const double *w, ulong n=0)
    {
        if ( 0!=n )  n_ = n;
        else  // zero terminated array w[]
        {
            n_ = 0;
            while ( w[n_]!=0 )  ++n_;
        }

        ctor_core();

        for (ulong i=0, j=n_-1;  i<n_;  ++i, --j)
        {
            h_[i] = w[i];

            if ( !(i&1) )  g_[j] = -h_[i];  // even indices
            else           g_[j] = +h_[i];  // odd indices
        }
    }

  [--snip--]
```

The wavelet conditions can be checked via

```
    bool check(double eps=1e-6) const
    {
        if ( fabs(norm_sqr(0)-1.0) > eps )  return false;

        for (ulong i=1;  i<n_/2;  ++i)
            if ( fabs(norm_sqr(i)) > eps )  return false;

        return true;
    }
```

where `norm_sqr()` computes the sums in the relations 26.1-5a and 26.1-5b:

```
    static double norm_sqr(const double *h, ulong n, ulong s=0)
    {
        s *= 2;  // Note!
        if ( s>=n )  return 0.0;

        double v = 0;
        for (ulong k=0,j=s;  j<n;  ++k,++j)  v += (h[k]*h[j]);
        return  v;
    }

    double norm_sqr(ulong s=0)  const  { return norm_sqr(h_, n_, s); }
```

A wavelet step can be implemented as [FXT: wavelet/wavelet.cc]:

```
void
wavelet_step(double *f, ulong n, const wavelet_filter &wf, double *t)
{
    const ulong nh = (n>>1);
    const ulong m = n-1;  // mask to compute modulo n (n is a power of two)
    for (ulong i=0,j=0; i<n; i+=2,++j) // i \in [0,2,4,..,n-2];  j \in [0,1,2,..,n/2-1]
    {
        double s = 0.0,  d = 0.0;
        for (ulong k=0; k<wf.n_; ++k)
        {
```

```
11              ulong w = (i+k) & m;
12              s += (wf.h_[k] * f[w]);
13              d += (wf.g_[k] * f[w]);
14          }
15          t[j] = s;
16          t[nh+j] = d;
17      }
18      copy(t, f, n);  // f[] := t[]
19  }
```

The wavelet transform itself is

```
1   void
2   wavelet(double *f, ulong ldn, const wavelet_filter &wf, ulong minm/*=2*/)
3   {
4       ulong n = (1UL<<ldn);
5       ALLOCA(double, t, n);
6       for (ulong m=n; m>=minm; m>>=1)  wavelet_step(f, m, wf, t);
7   }
```

The step for the inverse transform is [FXT: wavelet/invwavelet.cc]:

```
1   void
2   inverse_wavelet_step(double *f, ulong n, const wavelet_filter &wf, double *t)
3   {
4       const ulong nh = (n>>1);
5       const ulong m = n-1;  // mask to compute modulo n (n is a power of two)
6       null(t, n);  // t[] := [0,0,...,0]
7       for (ulong i=0, j=0;  i<n;  i+=2, ++j)
8       {
9           const double x = f[j],  y = f[nh+j];
10          for (ulong k=0;  k<wf.n_;  ++k)
11          {
12              ulong w = (i+k) & m;
13              t[w] += (wf.h_[k] * x);
14              t[w] += (wf.g_[k] * y);
15          }
16      }
17      copy(t, f, n);  // f[] := t[]
18  }
```

The inverse transform itself now is

```
1   void
2   inverse_wavelet(double *f, ulong ldn, const wavelet_filter &wf, ulong minm/*=2*/)
3   {
4       ulong n = (1UL<<ldn);
5       ALLOCA(double, t, n);
6       for (ulong m=minm; m<=n; m<<=1)  inverse_wavelet_step(f, m, wf, t);
7   }
```

A readable source about wavelets is [274].

26.3 Moment conditions

As the wavelet conditions do not uniquely define the wavelet filters on can impose additional properties for the filters used. We require that, for an $2n$-tap wavelet filter, the first $n/2$ moments vanish:

$$\sum_j (-1)^j\, h_j \;=\; 0 \tag{26.3-1a}$$

$$\sum_j (-j)^k\, h_j \;=\; 0 \quad \text{where} \quad k = 1, 2, 3, \ldots, n/2-1 \tag{26.3-1b}$$

One motivation for these *moment conditions* is that for reasonably smooth signals (for which a polynomial approximation is good) the transform coefficients from the high-pass filter (the d_k) will be close to zero. With compression schemes that simply discard transform coefficients with small values this is a desirable property.

The class [FXT: `class wavelet_filter` in wavelet/waveletfilter.h] has a method to compute the moments of the filter:

```
1       static double moment(const double *h, ulong n, ulong x=0)
2       {
3           if ( 0==x )
4           {
5               double v = 0.0;
6               for (ulong k=0; k<n; k+=2)  v += h[k];
7               for (ulong k=1; k<n; k+=2)  v -= h[k];
8               return v;
9           }
10
11          double dk;
12
13          double ve = 0;
14          dk = 2.0;
15          for (ulong k=2;  k<n;  k+=2, dk+=2.0)  ve += (pow(dk,x) * h[k]);
16
17          double vo = 0;
18          dk = 1.0;
19          for (ulong k=1;  k<n;  k+=2, dk+=2.0)  vo += (pow(dk,x) * h[k]);
20          return  ve - vo;
21      }
22
23      double moment(ulong x=0)  const  { return moment(h_, n_, x); }
```

Filter coefficients that satisfy the moment conditions are given in [FXT: wavelet/daubechies.cc]:

```
1       extern const double Daub1[] = {
2       +7.07106781186547524400844362104 8e-01,
3       +7.07106781186547524400844362104 8e-01 };
4
5       extern const double Daub2[] = {
6       +4.82962913144534143374871599864 4e-01,
7       +8.36516303737807905575293780916 8e-01,
8       +2.24143868042013381025972762240 4e-01,
9       -1.29409522551260381174449418812 0e-01 };
10
11      extern const double Daub3[] = {
12      +3.32670552950082615998511589139 0e-01,
13      +8.06891509311092576494493604088 7e-01,
14      +4.59877502118491570095151942147 6e-01,
15      -1.35011020010254588696389906699 3e-01,
16      -8.54412738820266616928191691817 7e-02,
17      +3.52262918857095366027406647155 1e-02 };
18
19      extern const double Daub4[] = {
20      +2.30377813308896500863291183044 0e-01,
21      +7.14846570552915647089921955273 9e-01,
22      +6.30880767929858907881716338300 6e-01,
23      -2.79837694168598542114137471800 7e-02,
24      -1.87034811719093084079570672789 0e-01,
25      +3.08413818355607636272193625349 5e-02,
26      +3.28830116668851997354075135492 4e-02,
27      -1.05974017850690321048832085240 2e-02 };
28
29        [--snip--]
30      extern const double Daub38[] = {...}
```

The names reflect the number $n/2$ of vanishing moments. Reversing or negating the sequence of filter coefficients leads to trivial variants that also satisfies the moment conditions.

For the filters of length $n \geq 6$ there are solutions that are essentially different. For $n = 6$ there is one complex solution besides `Daub3[]`:

```
1       -0.09556007476957763 + 0.0508627772544*I
2       +0.08121662052705924 + 0.1525883317632*I
3       +0.72145023542906591 + 0.1017255545088*I
4       +0.72145023542906591 - 0.1017255545088*I
5       +0.08121662052705924 - 0.1525883317632*I
6       -0.09556007476957763 - 0.0508627772544*I
```

For $n = 8$ there is, besides `Daub4[]`, an additional real solution (left), and a complex one (right):

```
1       -0.07576571478950221                +0.02152475910155493 + 0.0184283603930*I
2       -0.02963552764600249                -0.06571356411493559 + 0.0176790547520*I
3       +0.49761866763277498                -0.19397617446078878 - 0.1319957453155*I
4       +0.80373875180513208                +0.24627664139071534 - 0.2801719341011*I
5       +0.29785779560530605                +0.85723045931761476 - 0.0921418019654*I
6       -0.09921954357663353                +0.59199318785735184 + 0.2064584925288*I
7       -0.01260396726203130                +0.02232773722816661 + 0.2057091868878*I
8       +0.03222310060405146                -0.06544948394658407 + 0.0560343868202*I
```

The numbers of solutions grows exponentially with n (the minimal polynomial of any tap value has degree 2^n). The filters given in [FXT: wavelet/daubechies.cc] are the filters for the so-called *Daubechies wavelets* (some closed form expressions for the filter coefficients are given in [87]).

Filter coefficients that satisfy the wavelet and the moment conditions can be found by a Newton iteration for zeros of the function $F : \mathbb{R}^n \to \mathbb{R}^n$, $F(\vec{h}) := \vec{w}$ where $w_i = F_i(\vec{h}) = F_i(h_0, h_1, \ldots, h_5)$. For example, with $n = 6$ the F_i are defined by

```
1     F[1]:  h0^2 + h1^2 + h2^2 + h3^2 + h4^2 + h5^2 - 1
2     F[2]:  h2*h0 + h3*h1 + h4*h2 + h5*h3
3     F[3]:  h4*h0 + h5*h1
4     F[4]:  -h0 + h1 + -h2 + h3 + -h4 + h5
5     F[5]:  h1 + -2*h2 + 3*h3 + -4*h4 + 5*h5
6     F[6]:  h1 + -4*h2 + 9*h3 + -16*h4 + 25*h5
```

The derivative is given by the *Jacobi matrix* J. It has the components $J_{r,c} := \dfrac{dF_r}{dh_c}$. Its rows are

```
1     J[1]= [2*h0, 2*h1, 2*h2, 2*h3, 2*h4, 2*h5]
2     J[2]= [h2, h3, h0 + h4, h1 + h5, h2, h3]
3     J[3]= [h4, h5, 0, 0, h0, h1]
4     J[4]= [-1, 1, -1, 1, -1, 1]
5     J[5]= [0, 1, -2, 3, -4, 5]
6     J[6]= [0, 1, -4, 9, -16, 25]
```

Now iterate (the equivalent to Newton's iteration, $x_{k+1} := x_k - f(x_k)/f'(x_k)$)

$$\vec{h}_{k+1} \quad := \quad \vec{h}_k - J^{-1}(\vec{h}_k)\, F(\vec{h}_k) \tag{26.3-2}$$

The computations have to be carried out with a rather great precision to avoid loss of accuracy.

Part IV

Fast arithmetic

Chapter 27

Fast multiplication and exponentiation

The usual scheme for multiplication needs proportional N^2 operations for the multiplication of two N-digit numbers. This chapter describes multiplication algorithms that are asymptotically better than this, the Karatsuba algorithm, the Toom-Cook algorithms, and multiplication via FFTs. In addition, the left-to-right and right-to-left schemes for binary exponentiation are described.

27.1 Asymptotics of algorithms

An important feature of an algorithm is the number of operations that must be performed for the completion of a task of a certain size. For high precision computations one will take N as the length of the numbers counted in decimal digits or bits. For computations with square matrices one may take for N the number of rows, or the number of entries in the matrix. An operation is typically a (machine word) multiplication plus an addition, one could also simply count machine instructions.

We now consider the computational cost of a given algorithm. An algorithm is said to have some asymptotics $f(N)$ if it needs proportional $f(N)$ operations for a task of size N. We express proportionality with $f(N)$ as $\sim f(N)$. Some examples:

- Addition of two N-digit numbers needs $\sim N$ operations (here: machine word additions).
- Ordinary multiplication of two N-digit numbers needs $\sim N^2$ operations.
- Counting elements equal to a given value in an unsorted array of length N costs $\sim N$ operations. The operations are reads from memory and comparisons.
- Deciding whether a sorted array of length N contains a given value costs $\sim \log(N)$ operations when binary search is used.
- Computing the Fourier transform of a length-N sequence by definition, that is, as N sums each of length N, costs $\sim N^2$.
- The FFT algorithms compute the Fourier transform of a length-N sequence in $\sim N \log(N)$ operations
- Matrix multiplication (of two $N \times N$ matrices by the obvious algorithm) is $\sim N^3$ (N^2 sums, each of N products).
- When the problem size M for matrix multiplication is taken to be the number of elements of the matrix ($M = N^2$), then the cost is 'only' $\sim M^{3/2}$.

- Deciding whether a N-digit binary number is even or odd costs 1 operation (lookup whether the lowest bit is zero). The cost is independent of the problem size.

We have simplified the considerations by assuming that for a given algorithm $\sim f(N)$ there is a constant C so that

$$\lim_{N\to\infty} \frac{\text{work}(N)}{f(N)} \quad = \quad C \tag{27.1-1}$$

where work(N) is the actual cost for problem size N. The approximate cost of an algorithm is expressed by the symbol $\approx C \cdot f(N)$. The constant C is often referred to as the *hidden constant*.

The algorithm with the 'best' asymptotics wins for some, possibly huge, problem size N. For smaller N another algorithm will be superior. For the exact break-even point the hidden constants are, of course, important. For example, let the algorithm `mult1` take $\approx 1.0 \cdot N^2$ operations ($C = 1.0$), `mult2` take $\approx 8.0 \cdot N \log_2(N)$ operations ($C = 8.0$). Then for $N < 64$ the algorithm `mult1` is faster, while for $N > 64$ the algorithm `mult2` is faster. Completely different algorithms may be optimal for the same task at different problem sizes. The hidden constants can be so large that the asymptotically better algorithm is never practical for any feasible problem size.

With many algorithms it is only possible to give lower and upper (asymptotic) bounds, especially if the program flow depends on the processed data. Upper bounds are usually denoted with $O(f(N))$, for example, FFT multiplication is $O(N \log(N))$ and computations that take a constant number of operations independent of the problem size are $O(1)$.

The space requirements in all methods given is a constant times the problem size N. In-place algorithms (like the FFT) use space $\sim N$, the remaining cases (like radix sort) typically need $\sim 2\,N$. Note that it is possible to construct algorithms where the computational cost alone is not of much value: integer factorization (up to a certain maximal value) costs just one lookup if a precomputed table of factorizations is stored in an (insanely) big array.

For a more fine-grained approach to measuring the costs of algorithms see [95].

27.2 Splitting schemes for multiplication

Ordinary multiplication is $\sim N^2$. Assuming the hidden constant equals one the computation of the product of two million-digit numbers would require $\approx 10^{12}$ operations. On a machine that does 1 billion operations per second the multiplication would need 1000 seconds. The following schemes leads to algorithms with superior asymptotics.

27.2.1 2-way splitting: the Karatsuba algorithm

The following algorithm is due to A. Karatsuba and Y. Ofman, it was given 1963 in [160].

Split the numbers A and B (assumed to have approximately the same length) into two pieces

$$\begin{aligned} A &= x\,A_1 + A_0 \\ B &= x\,B_1 + B_0 \end{aligned} \tag{27.2-1}$$

where x is a power of the radix (for decimal numbers the radix is 10) close to the half length of A and B. The usual multiplication scheme needs 4 multiplications with half precision for one multiplication with full precision:

$$A\,B \quad = \quad A_0 \cdot B_0 + x\,(A_0 \cdot B_1 + B_0 \cdot A_1) + x^2\,A_1 \cdot B_1 \tag{27.2-2}$$

Only the multiplications $A_i \cdot B_j$ need to be considered. The multiplications by x, a power of the radix, are only shifts. If we use the relation

$$A\,B \;=\; (1+x)\,A_0 \cdot B_0 + x\,(A_1 - A_0) \cdot (B_0 - B_1) + (x + x^2)\,A_1 \cdot B_1 \tag{27.2-3}$$

we need 3 multiplications with half precision for one multiplication with full precision. Applying the scheme recursively until the numbers to multiply are of machine size we obtain an algorithm whose asymptotic cost is $\sim N^{\log_2(3)} \approx N^{1.585}$. An alternative form of relation 27.2-3 is

$$A\,B \;=\; (1-x)\,A_0 \cdot B_0 + x\,(A_1 + A_0) \cdot (B_0 + B_1) + (x^2 - x)\,A_1 \cdot B_1 \tag{27.2-4}$$

For squaring use either of the following schemes

$$A^2 \;=\; (1+x)\,A_0^2 - x\,(A_1 - A_0)^2 + (x + x^2)\,A_1^2 \tag{27.2-5a}$$
$$A^2 \;=\; (1-x)\,A_0^2 + x\,(A_1 + A_0)^2 + (x^2 - x)\,A_1^2 \tag{27.2-5b}$$

We compute $8231^2 = 67749361$ with the first relation (27.2-5a):

```
8231^2  ==  (100*82+31)^2
        ==  (1+100)*31^2  - 100*(82-31)^2 + (100+100^2)*82^2
        ==  (1+100)*[961] - 100*[2601]    + (100+100^2)*[6724]
        ==  961 + 96100   - 260100        + 672400 + 67240000
        ==  67749361
```

Assume that the hidden constant equals 2 as there is more bookkeeping overhead than with the usual algorithm. Computing the product of two million-digit numbers would require $\approx 2 \cdot (10^6)^{1.585} \approx 6.47 \cdot 10^9$ operations, taking about 6.5 seconds on our computer.

The Karatsuba scheme for polynomial multiplication is given in section 38.2 on page 821.

27.2.2 3-way splitting

One can extend the above idea by splitting U and V into more than two pieces each, the resulting method is called *Toom-Cook algorithm* (the method is called *Toom algorithm* in [94], and *Cook-Toom algorithm* in [2]).

27.2.2.1 Zimmermann's 3-way multiplication

```
A = a2*x^2 + a1*x + a0
B = b2*x^2 + b1*x + b0

S0 = a0 * b0
S1 = (a2+a1+a0) * (b2+b1+b0)
S2 = (4*a2+2*a1+a0) * (4*b2+2*b1+b0)
S3 = (a2-a1+a0) * (b2-b1+b0)
S4 = a2 * b2

T1 = 2*S3 + S2
T1 /= 3  \\ division by 3
T1 += S0
T1 /= 2
T1 -= 2*S4
T2 = (S1 + S3)/2
S1 -= T1
S2 = T2 - S0 - S4
S3 = T1 - T2

P = S4*x^4 + S3*x^3 + S2*x^2 + S1*x + S0
P - A*B  \\ == zero
```

Figure 27.2-A: Implementation of Zimmermann's 3-way multiplication scheme in pari/gp.

The good scheme for 3-way splitting is due to Paul Zimmermann. We compute the product $C = A \cdot B$ of two numbers, A and B

$$A = a_2\,x^2 + a_1\,x + a_0 \tag{27.2-6a}$$
$$B = b_2\,x^2 + b_1\,x + b_0 \tag{27.2-6b}$$
$$C = A \cdot B = c_4\,x^4 + c_3\,x^3 + c_2\,x^2 + c_1\,x + c_0 \tag{27.2-6c}$$

by the following scheme (taken from [86]): set

$$S_0 := a_0 \cdot b_0 \tag{27.2-6d}$$
$$S_1 := (a_2 + a_1 + a_0) \cdot (b_2 + b_1 + b_0) \tag{27.2-6e}$$
$$S_2 := (4\,a_2 + 2\,a_1 + a_0) \cdot (4\,b_2 + 2\,b_1 + b_0) \tag{27.2-6f}$$
$$S_3 := (a_2 - a_1 + a_0) \cdot (b_2 - b_1 + b_0) \tag{27.2-6g}$$
$$S_4 := a_2 \cdot b_2 \tag{27.2-6h}$$

This costs 5 multiplications of length $N/3$. We already have found $c_0 = S_0$ and $c_4 = S_4$. We determine c_1, c_2, and c_3 by the following assignments (in the given order):

$$T_1 := 2\,S_3 + S_2 \qquad (= 18\,c_4 + 6\,c_3 + 6\,c_2 + 3\,c_0) \tag{27.2-6i}$$
$$T_1 := T_1/3 \qquad (= 6\,c_4 + 2\,c_3 + 2\,c_2 + c_0) \qquad \text{exact division by 3} \tag{27.2-6j}$$
$$T_1 := T_1 + S_0 \qquad (= 6\,c_4 + 2\,c_3 + 2\,c_2 + 2\,c_0) \tag{27.2-6k}$$
$$T_1 := T_1/2 \qquad (= 3\,c_4 + c_3 + c_2 + c_0) \tag{27.2-6l}$$
$$T_1 := T_1 - 2\,S_4 \qquad (= c_4 + c_3 + c_2 + c_0) \tag{27.2-6m}$$
$$T_2 := (S_1 + S_3)/2 \qquad (= c_4 + c_2 + c_0) \tag{27.2-6n}$$
$$S_1 := S_1 - T_1 \qquad (= c_1) \qquad \text{wrong in cited paper} \tag{27.2-6o}$$
$$S_2 := T_2 - S_0 - S_4 \qquad (= c_2) \tag{27.2-6p}$$
$$S_3 := T_1 - T_2 \qquad (= c_3) \tag{27.2-6q}$$

Now we have

$$C = A \cdot B = S_4\,x^4 + S_3\,x^3 + S_2\,x^2 + S_1\,x + S_0 \tag{27.2-6r}$$

The complexity of recursive multiplication based on this splitting scheme is $N^{\log_3(5)} \approx N^{1.465}$. Assume that the hidden constant again equals 2. Then the computation of the product of two million-digit numbers would require $\approx 2 \cdot (10^6)^{1.465} \approx 1.23 \cdot 10^9$ operations, taking about 1.2 seconds on our computer.

Note the division by 3 in relation 27.2-6j. A division by a constant (that is not a power of two) cannot be avoided in n-way splitting schemes for multiplication for $n \geq 3$. There are squaring schemes that do not involve such divisions.

27.2.2.2 3-way multiplication by Bodrato and Zanoni

An alternative algorithm for 3-way splitting is suggested in [50]: setup S_0, S_1, ..., S_4 as in relations 27.2-6d... 27.2-6h, then compute, in the given order,

$$S_2 := (S_2 - S_3)/3 \qquad (= 5\,c_4 + 3\,c_3 + c_2 + c_1) \qquad \text{exact division by 3} \tag{27.2-7a}$$
$$S_3 := (S_1 - S_3)/2 \qquad (= c_3 + c_1) \tag{27.2-7b}$$
$$S_1 := S_1 - S_0 \qquad (= c_4 + c_3 + c_2 + c_1) \tag{27.2-7c}$$
$$S_2 := (S_2 - S_1)/2 \qquad (= 2\,c_4 + c_3) \tag{27.2-7d}$$
$$S_1 := S_1 - S_3 - S_4 \qquad (= c_2) \tag{27.2-7e}$$
$$S_2 := S_2 - 2\,S_4 \qquad (= c_3) \tag{27.2-7f}$$
$$S_3 := S_3 - S_2 \qquad (= c_1) \tag{27.2-7g}$$

```
A = a2*x^2 + a1*x + a0
B = b2*x^2 + b1*x + b0

S0 = a0 * b0
S1 = (a2+a1+a0) * (b2+b1+b0)
S2 = (4*a2+2*a1+a0) * (4*b2+2*b1+b0)
S3 = (a2-a1+a0) * (b2-b1+b0)
S4 = a2 * b2

S2 = (S2 - S3)/3  \\ division by 3
S3 = (S1 - S3)/2
S1 = S1 - S0
S2 = (S2 - S1)/2
S1 = S1 - S3 - S4
S2 = S2 - 2*S4
S3 = S3 - S2

P = S4*x^4+ S2*x^3+ S1*x^2+ S3*x + S0

P - A*B \\ == zero
```

Figure 27.2-B: Implementation of the 3-way multiplication scheme of Bodrato and Zanoni.

Now we have (note the order of the coefficients S_i)

$$C \;=\; A \cdot B \;=\; S_4\,x^4 + S_2\,x^3 + S_1\,x^2 + S_3\,x + S_0 \tag{27.2-7h}$$

The scheme requires only one multiplication by two, Zimmermann's scheme involves two.

27.2.2.3 3-way squaring

The following scheme is taken from [86]. To compute the square $C = A^2$ of a number A

$$A \;=\; a_2\,x^2 + a_1\,x + a_0 \tag{27.2-8a}$$
$$C \;=\; A^2 \;=\; S_4\,x^4 + S_3\,x^3 + S_2\,x^2 + S_1\,x + S_0 \tag{27.2-8b}$$

set

$$S_0 \;:=\; a_0^2 \tag{27.2-8c}$$
$$S_1 \;:=\; (a_2 + a_1 + a_0)^2 \tag{27.2-8d}$$
$$S_2 \;:=\; (a_2 - a_1 + a_0)^2 \tag{27.2-8e}$$
$$S_3 \;:=\; 2\,a_1 \cdot a_2 \tag{27.2-8f}$$
$$S_4 \;:=\; a_2^2 \tag{27.2-8g}$$
$$\tag{27.2-8h}$$

This costs 4 squarings and 1 multiplication of length $N/3$. The quantities S_0, S_3, and S_4 are already correct. Determine S_1 and S_2 via

$$T_1 \;:=\; (S_1 + S_2)/2 \tag{27.2-8i}$$
$$S_1 \;:=\; S_1 - T_1 - S_3 \tag{27.2-8j}$$
$$S_2 \;:=\; T_1 - S_4 - S_0 \tag{27.2-8k}$$

27.2.3 4-way splitting

27.2.3.1 4-way multiplication

An elegant and clean scheme for 4-way splitting of a multiplication is given by Bodrato and Zanoni in [51]. A pari/gp implementation is shown in figure 27.2-C. The algorithm has asymptotics $\sim O(n^{\log_4(7)}) \approx O(n^{1.403})$. In general, an s-way splitting scheme will be $\sim O(n^{\log_s(2s+1)})$.

```
A = a3*x^3 + a2*x^2 + a1*x + a0
B = b3*x^3 + b2*x^2 + b1*x + b0

S1 = a3*b3
S2 = (8*a3+4*a2+2*a1+a0)*(8*b3+4*b2+2*b1+b0)
S3 = (+a3+a2+a1+a0)*(+b3+b2+b1+b0)
S4 = (-a3+a2-a1+a0)*(-b3+b2-b1+b0)
S5 = (+8*a0+4*a1+2*a2+a3)*(+8*b0+4*b1+2*b2+b3);
S6 = (-8*a0+4*a1-2*a2+a3)*(-8*b0+4*b1-2*b2+b3)
S7 = a0*b0


S2 += S5
S4 -= S3
S6 -= S5
S4 /= 2  \\
S5 -= S1
S5 -= (64*S7)
S3 += S4
S5 *= 2;  S5 += S6

S2 -= (65*S3)
S3 -= S1
S3 -= S7
S4 = -S4 \\
S6 = -S6 \\
S2 += (45*S3)
S5 -= (8*S3)
S5 /= 24  \\ division by 24

S6 -= S2
S2 -= (16*S4)
S2 /= 18  \\ division by 18
S3 -= S5
S4 -= S2
S6 += (30*S2)
S6 /= 60   \\ division by 60
S2 -= S6

P = S1*x^6 + S2*x^5 + S3*x^4 + S4*x^3 + S5*x^2 + S6*x + S7;

P - A*B \\ == zero
```

Figure 27.2-C: Implementation of the 4-way multiplication scheme in pari/gp.

27.2.3.2 4-way squaring

The following scheme is taken from [86]

$$A = a_3\,x^3 + a_2\,x^2 + a_1\,x + a_0 \tag{27.2-9a}$$
$$C = A^2 = c_6\,x^6 + c_5\,x^5 + c_4\,x^4 + c_3\,x^3 + c_2\,x^2 + c_1\,x + c_0 \tag{27.2-9b}$$

Set

$$S_1 := a_0^2 \tag{27.2-9c}$$
$$S_2 := 2\,a_0 \cdot a_1 \tag{27.2-9d}$$
$$S_3 := (a_0 + a_1 - a_2 - a_3) \cdot (a_0 - a_1 - a_2 + a_3) \tag{27.2-9e}$$
$$S_4 := (a_0 + a_1 + a_2 + a_3)^2 \tag{27.2-9f}$$
$$S_5 := 2\,(a_0 - a_2) \cdot (a_1 - a_3) \tag{27.2-9g}$$
$$S_6 := 2\,a_3 \cdot a_2 \tag{27.2-9h}$$
$$S_7 := a_3^2 \tag{27.2-9i}$$

```
A = a3*x^3 + a2*x^2 + a1*x + a0

S1 = a0^2
S2 = 2 *a0 *a1
S3 = (a0 + a1 - a2 - a3) * (a0 - a1 - a2 + a3)
S4 = (a0 + a1 + a2 + a3 )^2
S5 = 2*(a0 - a2)*(a1 - a3)
S6 = 2*a3*a2
S7 = a3^2

T1 = S3 + S4
T2 = (T1 + S5)/2
T3 = S2 + S6
T4 = T2 - T3
T5 = T3 - S5
T6 = T4 - S3
T7 = T4 - S1
T8 = T6 - S7

P = S7 *x^6 + S6 *x^5 + T7 *x^4 + T5 *x^3 + T8 *x^2 + S2 *x + S1
P - A^2  \\ == zero
```

Figure 27.2-D: Implementation of the 4-way squaring scheme in pari/gp.

Then set, in the given order,

$$\begin{aligned} T_1 &:= S_3 + S_4 && \text{(27.2-9j)}\\ T_2 &:= (T_1 + S_5)/2 && \text{(27.2-9k)}\\ T_3 &:= S_2 + S_6 && \text{(27.2-9l)}\\ T_4 &:= T_2 - T_3 && \text{(27.2-9m)}\\ T_5 &:= T_3 - S_5 && \text{(27.2-9n)}\\ T_6 &:= T_4 - S_3 && \text{(27.2-9o)}\\ T_7 &:= T_4 - S_1 && \text{(27.2-9p)}\\ T_8 &:= T_6 - S_7 && \text{(27.2-9q)} \end{aligned}$$

The square then equals

$$C = S_7\,x^6 + S_6\,x^5 + T_7\,x^4 + T_5\,x^3 + T_8\,x^2 + S_2\,x + S_1 \qquad \text{(27.2-9r)}$$

27.2.4 5-way splitting

27.2.4.1 5-way multiplication

The scheme for 5-way splitting of a multiplication shown in figure 27.2-E is given in [51]. As for the 4-way multiplication scheme, no temporaries are used.

27.2.4.2 5-way squaring

We describe the 5-way squaring scheme given in [50]. Let

$$\begin{aligned} A &= a_4\,x^4 + a_3\,x^3 + a_2\,x^2 + a_1\,x + a_0 && \text{(27.2-10)}\\ C &= A^2 = c_8\,x^8 + c_7\,x^7 + c_6\,x^6 + c_5\,x^5 + c_4\,x^4 + c_3\,x^3 + c_2\,x^2 + c_1\,x + c_0 && \text{(27.2-11)} \end{aligned}$$

```
A = a4*x^4 + a3*x^3 + a2*x^2 + a1*x + a0
B = b4*x^4 + b3*x^3 + b2*x^2 + b1*x + b0


S1 = a4*b4

S2 = (a0-2*a1+4*a2-8*a3+16*a4)*(b0-2*b1+4*b2-8*b3+16*b4)
S5 = (a0+2*a1+4*a2+8*a3+16*a4)*(b0+2*b1+4*b2+8*b3+16*b4)

S3 = (a4+2*a3+4*a2+8*a1+16*a0)*(b4+2*b3+4*b2+8*b1+16*b0)
S8 = (a4-2*a3+4*a2-8*a1+16*a0)*(b4-2*b3+4*b2-8*b1+16*b0)

S4 = (a0+4*a1+16*a2+64*a3+256*a4)*(b0+4*b1+16*b2+64*b3+256*b4)

S6 = (a0-a1+a2-a3+a4)*(b0-b1+b2-b3+b4)
S7 = (a0+a1+a2+a3+a4)*(b0+b1+b2+b3+b4)

S9 = a0*b0


S6 -= S7
S2 -= S5
S4 -= S9
S4 -= (2^16*S1)
S8 -= S3
S6 /= 2 \\
S5 *= 2;  S5 += S2
S2 = -S2 \\
S8 = -S8 \\
S7 += S6
S6 = -S6 \\
S3 -= S7
S5 -= (512*S7)
S3 *= 2;  S3 -= S8

S7 -= S1
S7 -= S9
S8 += S2
S5 += S3
S8 -= (80*S6)
S3 -= (510*S9)
S4 -= S2
S3 *= 3;  S3 += S5
S8 /= 180  \\ division by 180
S5 += (378*S7)
S2 /= 4
S6 -= S2
S5 /= (-72)  \\ division by -72
S3 /= (-360)  \\ division by -360

S2 -= S8
S7 -= S3
S4 -= (256*S5)
S3 -= S5
S4 -= (4096*S3)
S4 -= (16*S7)
S4 += (256*S6)
S6 += S2
S2 *= 180;  S2 += S4
S2 /= 11340  \\ division by 11340
S4 += (720*S6)
S4 /= (-2160)  \\ division by -2160
S6 -= S4
S8 -= S2


P = S1*x^8 + S2*x^7 + S3*x^6 + S4*x^5 + S5*x^4 + S6*x^3 + S7*x^2 + S8*x + S9;

P - A*B \\ == zero
```

Figure 27.2-E: Implementation of the 5-way multiplication scheme in pari/gp.

```
A = a4*x^4+ a3*x^3 + a2*x^2 + a1*x + a0

S1 = a0^2
S2 = a4^2
S3 = (a0 + a1 + a2 + a3 + a4)^2
S4 = (a0 - a1 + a2 - a3 + a4)^2
S5 = 2* (a0-a2+a4) * (a1-a3)
S6 = (a0 + a1 - a2 - a3 + a4) * (a0 - a1 - a2 + a3 + a4)
S7 = (a1 + a2 - a4) * (a1 - a2 - a4 + 2*(a0-a3))
S8 = 2*a0*a1
S9 = 2*a3*a4

S4 = (S4+S3)/2
S3 = S3-S4
S6 = (S6+S4)/2
S5 = (-S5+S3)/2
S4 = S4-S6
S3 = S3-S5-S8
S6 = S6-S2-S1
S5 = S5-S9
S7 = S7-S2-S8-S9+S6+S3
S4 = S4-S7

P = S2*x^8+S9*x^7+S4*x^6+S3*x^5+S6*x^4+S5*x^3+S7*x^2+S8*x+S1
P - A^2  \\ == zero
```

Figure 27.2-F: Implementation of the 5-way squaring scheme,

Set

$$
\begin{aligned}
S_1 &:= a_0^2 && \text{(27.2-12a)}\\
S_2 &:= a_4^2 && \text{(27.2-12b)}\\
S_3 &:= (a_0 + a_1 + a_2 + a_3 + a_4)^2 && \text{(27.2-12c)}\\
S_4 &:= (a_0 - a_1 + a_2 - a_3 + a_4)^2 && \text{(27.2-12d)}\\
S_5 &:= 2\,(a_0 - a_2 + a_4) \cdot (a_1 - a_3) && \text{(27.2-12e)}\\
S_6 &:= (a_0 + a_1 - a_2 - a_3 + a_4) \cdot (a_0 - a_1 - a_2 + a_3 + a_4) && \text{(27.2-12f)}\\
S_7 &:= (a_1 + a_2 - a_4) \cdot (a_1 - a_2 - a_4 + 2\,(a_0 - a_3)) && \text{(27.2-12g)}\\
S_8 &:= 2\,a_0 \cdot a_1 && \text{(27.2-12h)}\\
S_9 &:= 2\,a_3 \cdot a_4 && \text{(27.2-12i)}
\end{aligned}
$$

Further set, in the order given,

$$
\begin{aligned}
S_4 &= (S_4 + S_3)/2 \qquad (= c_0 + c_2 + c_4 + c_6 + c_8) && \text{(27.2-13a)}\\
S_3 &= S_3 - S_4 \qquad (= c_1 + c_3 + c_5 + c_7) && \text{(27.2-13b)}\\
S_6 &= (S_6 + S_4)/2 \qquad (= c_0 + c_4 + c_8) && \text{(27.2-13c)}\\
S_5 &= (-S_5 + S_3)/2 \qquad (= c_3 + c_7) && \text{(27.2-13d)}\\
S_4 &= S_4 - S_6 \qquad (= c_2 + c_6) && \text{(27.2-13e)}\\
S_3 &= S_3 - S_5 - S_8 \qquad (= c_5) && \text{(27.2-13f)}\\
S_6 &= S_6 - S_2 - S_1 \qquad (= c_4) && \text{(27.2-13g)}\\
S_5 &= S_5 - S_9 \qquad (= c_3) && \text{(27.2-13h)}\\
S_7 &= S_7 - S_2 - S_8 - S_9 + S_6 + S_3 \qquad (= c_2) && \text{(27.2-13i)}\\
S_4 &= S_4 - S_7 \qquad (= c_6) && \text{(27.2-13j)}
\end{aligned}
$$

Now we have (note the order of the coefficients S_i)

$$C \;=\; A^2 \;=\; S_2\,x^8 + S_9\,x^7 + S_4\,x^6 + S_3\,x^5 + S_6\,x^4 + S_5\,x^3 + S_7\,x^2 + S_8\,x + S_1 \qquad \text{(27.2-13k)}$$

```
A = a4*x^4+ a3*x^3 + a2*x^2 + a1*x + a0

S1 = a0^2
S2 = a4^2
[ ...S9,  as before]

T1 = S1 + 2*S2 - S7 + 2*S8 + S9
T2 = S3 - S4
T3 = 2*S5
T4 = T2 + T3
T5 = T2 - T3
T6 = T4/4
T7 = T5/4 - S9
T8 = T1 - T6 - S6
T9 = T6 - S8
T10 = S3 + S6
T11 = (T10 + S4 + S6)/4
T12 = T11 - S1 - S2
T13 = (T10 + S5)/2
T14 = T13 - T1

P = S2*x^8 + S9*x^7 + T8*x^6 + T9*x^5 +  T12*x^4 + T7*x^3 + T14*x^2 + S8*x + S1
P - A^2 \\ == zero
```

Figure 27.2-G: Implementation of the alternative 5-way squaring scheme in pari/gp. Definition of $S_1,\ldots,S_9$ as in figure 27.2-F.

27.2.4.3 Alternative 5-way squaring scheme

The following scheme is taken from [86], we correct some errors in the paper. Setup $S_1,\ldots,S_9$ as given by relations 27.2-12a...27.2-12i, then compute, in the given order,

$$\begin{aligned}
T_1 &:= S_1 + 2\,S_2 - S_7 + 2\,S_8 + S_9 && \text{(27.2-14a)}\\
T_2 &:= S_3 - S_4 && \text{(27.2-14b)}\\
T_3 &:= 2\,S_5 && \text{(27.2-14c)}\\
T_4 &:= T_2 + T_3 && \text{(27.2-14d)}\\
T_5 &:= T_2 - T_3 && \text{(27.2-14e)}\\
T_6 &:= T_4/4 && \text{(27.2-14f)}\\
T_7 &:= T_5/4 - S_9 && \text{(27.2-14g)}\\
T_8 &:= T_1 - T_6 - S_6 && \text{(27.2-14h)}\\
T_9 &:= T_6 - S_8 && \text{(27.2-14i)}\\
T_{10} &:= S_3 + S_6 && \text{(27.2-14j)}\\
T_{11} &:= (T_{10} + S_4 + S_6)/4 && \text{(27.2-14k)}\\
T_{12} &:= T_{11} - S_1 - S_2 \qquad \text{wrong in cited paper} && \text{(27.2-14l)}\\
T_{13} &:= (T_{10} + S_5)/2 && \text{(27.2-14m)}\\
T_{14} &:= T_{13} - T_1 && \text{(27.2-14n)}
\end{aligned}$$

We have (note that the coefficients for x^4 and x^2 are wrong in the cited paper):

$$C = S_2\,x^8 + S_9\,x^7 + T_8\,x^6 + T_9\,x^5 + T_{12}\,x^4 + T_7\,x^3 + T_{14}\,x^2 + S_8\,x + S_1 \qquad \text{(27.2-14o)}$$

27.3 Fast multiplication via FFT

Multiplication of two numbers is essentially a convolution of the sequences of their digits. The (linear) convolution of the two sequences $a_k, b_k, k = 0\ldots N-1$ is defined as the sequence c where

$$c_k := \sum_{i,j=0;\ i+j=k}^{N-1} a_i b_j \qquad k = 0\ldots 2N-2 \qquad \text{(27.3-1)}$$

A (n-digit) number written in radix R as

$$a_{n-1}\, a_{n-2}\, \ldots\, a_2\, a_1\, a_0 \tag{27.3-2}$$

denotes a quantity of

$$\sum_{i=0}^{n-1} a_i \cdot R^i \quad = \quad a_{n-1} \cdot R^{n-1} + a_{n-2} \cdot R^{n-2} + \ldots + a_1 \cdot R + a_0 \tag{27.3-3}$$

This means, the digits can be identified with coefficients of a polynomial in R. For example, with decimal numbers one has $R = 10$, and the number 578 equals $5 \cdot 10^2 + 7 \cdot 10^1 + 8 \cdot 10^0$. The product of two numbers is almost the polynomial product

$$\sum_{k=0}^{2N-2} c_k\, R^k \quad := \quad \sum_{i=0}^{N-1} a_i\, R^i \cdot \sum_{j=0}^{N-1} b_j\, R^j \tag{27.3-4}$$

The c_k are found by comparing coefficients. One easily checks that the c_k must satisfy the convolution equation 27.3-1. As the c_k can be greater than 'nine' (that is, $R - 1$), the result has to be 'fixed' using *carry* operations: Go from right to left, replace c_k by $c'_k = c_k \bmod R$ and add $(c_k - c'_k)/R$ to its left neighbor.

An example: usually one would multiply the numbers 82 and 34 as follows:

	82	×	34	
	3	^{3}2	8	
2	^{2}4	6		
= 2	7	8	8	

We have seen that the carries can be delayed to the end of the computation:

	82	×	34	
		32	8	
	24	6		
	24	38	8	
= 2	^{2} 7	^{3} 8	8	

... which is really polynomial multiplication (which in turn is a convolution of the coefficients):

	$(8\,x + 2)$	×	$(3\,x + 4)$	
		$32\,x$	8	
	$24\,x^2$	$6\,x$		
=	$24\,x^2$	$+38\,x$		$+8$

The value of the polynomial $24\,x^2 + 38\,x + 8$ for $x = 10$ is 2788.

Convolution can be done efficiently using the Fast Fourier Transform (FFT): convolution is a simple (element-wise) multiplication in Fourier space. The FFT itself takes $\sim N \cdot \log N$ operations. Instead of the direct convolution ($\sim N^2$) one proceeds as follows:

- Compute the FFTs of both factors.
- Multiply the transformed sequences element-wise.
- Compute inverse transform of the product.

To understand why this actually works note that (1) the multiplication of two polynomials can be achieved by the (more complicated) scheme:

- evaluate both polynomials at sufficiently many points (at least one more point than the degree of the product polynomial c: $\deg c = \deg a + \deg b$)

- element-wise multiply the values found
- find the polynomial corresponding to those (product-)values

and (2) that *the FFT is an algorithm for the parallel evaluation of a given polynomial at many points, namely the roots of unity.* (3) the inverse FFT is an algorithm to find (the coefficients of) a polynomial whose values are given at the roots of unity.

You might be surprised if you always thought of the FFT as an algorithm for the 'decomposition into frequencies'. There is no problem with either of these notions.

Re-launching our example ($82 \cdot 34 = 2788$), we use the fourth roots of unity ± 1 and $\pm i$:

	$a = (8\,x + 2)$	$\times$	$b = (3\,x + 4)$	$c = a\,b$
$+1$	$+10$		$+7$	$+70$
$+i$	$+8i + 2$		$+3i + 4$	$+38i - 16$
-1	-6		$+1$	-6
$-i$	$-8i + 2$		$-3i + 4$	$-38i - 16$
				$c = (24\,x^2 + 38\,x + 8)$

This table has to be read as follows: first the given polynomials a and b are evaluated at the points given in the left column, thereby the columns below a and b are filled. Then the values are multiplied to fill the column below c, giving the values of c at the points. Finally, the actual polynomial c is found from those values, resulting in the lower right entry. You may find it instructive to verify that a 4-point FFT really evaluates a, b by transforming the sequences 0, 0, 8, 2 and 0, 0, 3, 4 by hand. The backward transform of 70, $38i - 16$, -6, $-38i - 16$ should produce the final result given for c.

The operation count is dominated by that of the FFTs (the element-wise multiplication is of course $\sim N$), so the whole fast convolution algorithm takes $\sim N \cdot \log N$ operations. The following carry operation is also $\sim N$ and can therefore be neglected when counting operations.

Assume the hidden constant equals five. Multiplying our million-digit numbers will need about

$$5 \cdot 10^6 \log_2(10^6) \quad \approx \quad 5 \cdot 10^6 \cdot 20 \;=\; 10^8 \;=\; 0.1 \cdot 10^9$$

operations, taking approximately a tenth second on our computer.

Strictly speaking $N \cdot \log N$ is not really the truth: it has to be $N \cdot \log N \cdot \log\log N$. This is because the sums in the convolutions have to be represented as exact integers. The biggest term C that can possibly occur is approximately $N\,R^2$ for a number with N digits (see next section). Therefore, working with some fixed radix R one has to compute the FFTs with $\log N$ bits precision, leading to an operation count of $N \cdot \log N \cdot \log N$. The slightly better $N \cdot \log N \cdot \log\log N$ can be obtained by recursive use of FFT multiplies. For realistic applications (where the sums in the convolution all fit into the machine type floating point numbers) it is safe to think of FFT multiplication being proportional $N \cdot \log N$.

Several alternative multiplication algorithms are given in [168, ch.4.3.3]. See [176] on how far the idea "polynomials for numbers" can be carried and where it fails.

27.4 Radix/precision considerations with FFT multiplication

This section describes the dependencies between the radix of the number and the achievable precision when using FFT multiplication.

We use (unsigned) 16-bit words for the digits. Thereby the radix of the numbers can be in the range 2, 3, ..., 65536($= 2^{16}$). When working in base ten one will actually use 'super-digits' of base 10,000, the largest power of ten that fits into a 16-bit word. These super-digits are called `LIMB`s in hfloat.

With very large precision one cannot always use the greatest power of the desired base. This is due to the fact that the components of the convolution must be representable as integer numbers with the data type

used for the FFTs: The cumulative sums c_k have to be represented precisely enough to distinguish every (integer) quantity from the next bigger (or smaller) value. The highest possible value for a c_k will appear in the middle of the product and when multiplicand and multiplier consist of 'nines' (that is $R-1$) only. For radix R and a precision of N `LIMB`s Let the maximal possible value be C, then

$$C \quad = \quad N\,(R-1)^2 \tag{27.4-1}$$

Note that with FFT based convolution the *absolute value* of the central term can in fact equal $|C| = N^2\,(R-1)^2$. But there is no need to distinguish that many integers. After dividing by N we are back at relation 27.4-1.

The number of bits to represent C exactly is the integer greater or equal to

$$\log_2(N\,(R-1)^2) \quad = \quad \log_2 N + 2\,\log_2(R-1) \tag{27.4-2}$$

Due to numerical errors there must be a few more bits for safety. If computations are made using double-precision floating point numbers (C-type `double`) one typically has a mantissa of 53 bits[1] then we need to have

$$M \quad \geq \quad \log_2 N + 2\,\log_2(R-1) + S \tag{27.4-3}$$

where M :=mantissa-bits and S :=safety-bits. Using $\log_2(R-1) < \log_2(R)$:

$$N_{max}(R) \quad = \quad 2^{M-S-2\,\log_2(R)} \tag{27.4-4}$$

Suppose we have $M = 53$ mantissa-bits and require $S = 3$ safety-bits. With base 2 numbers one could use radix $R = 2^{16}$ for precisions up to a length of $N_{max} = 2^{53-3-2\cdot 16} = 256k$ `LIMB`s. Corresponding are 4096 kilo bits and = 1024 kilo hex digits. For greater lengths smaller radices have to be used according to the following table (extra horizontal line at the 16 bit limit for `LIMB`s):

Radix R	max # `LIMB`s	max # hex digits	max # bits
$2^{10} = 1024$	$1048,576\,k$	$2621,440\,k$	$10240\,M$
$2^{11} = 2048$	$262,144\,k$	$720,896\,k$	$2816\,M$
$2^{12} = 4096$	$65,536\,k$	$196,608\,k$	$768\,M$
$2^{13} = 8192$	$16384\,k$	$53,248\,k$	$208\,M$
$2^{14} = 16384$	$4096\,k$	$14,336\,k$	$56\,M$
$2^{15} = 32768$	$1024\,k$	$3840\,k$	$15\,M$
$2^{16} = 65536$	$256\,k$	$1024\,k$	$4\,M$
$2^{17} = 128\,k$	$64\,k$	$272\,k$	$1062\,k$
$2^{18} = 256\,k$	$16\,k$	$72\,k$	$281\,k$
$2^{19} = 512\,k$	$4\,k$	$19\,k$	$74\,k$
$2^{20} = 1\,M$	$1\,k$	$5\,k$	$19\,k$
$2^{21} = 2\,M$	256	1300	5120

For decimal numbers:

Radix R	max # `LIMB`s	max # digits	max # bits
10^2	$110\,G$	$220\,G$	$730\,G$
10^3	$1100\,M$	$3300\,M$	$11\,G$
10^4	$11\,M$	$44\,M$	$146\,M$
10^5	$110\,k$	$550\,k$	$1826\,k$
10^6	$1\,k$	$6,597$	$22\,k$
10^7	11	77	255

Summarizing:

- For decimal digits and precisions up to 11 million `LIMB`s use radix 10,000. (corresponding to more than 44 million decimal digits), for even greater precisions choose radix 1,000.

[1] Of which only the 52 least significant bits are physically present, the most significant bit is implied to be always set.

- For hexadecimal digits and precisions up to 256,000 LIMBs use radix 65,536 (corresponding to more than 1 million hexadecimal digits), for even greater precisions choose radix 4,096.

27.5 The sum-of-digits test

With high-precision calculations it is mandatory to add a sanity check to the multiplication routines. That is, false results due to loss of accuracy should (with high probability) be detected via the *sum-of-digits test*:

When computing the product $c = a \cdot b$ with radix-R numbers then compute the values ('sums of digits') $s_a = a \bmod (R-1)$ and $s_b = b \bmod (R-1)$ after the multiplication compute $s_c = c \bmod (R-1)$ and compare s_c to $s_m = s_a \cdot s_b \bmod (R-1)$. If $s_c \neq s_m$ then an error has occurred in the computation of c.

The sum-of-digits function s_a can for a radix-R, length-n number a be computed as [FXT: mult/auxil.cc]:

```
1   ulong
2   sum_of_digits(const LIMB *a, ulong n, ulong nine, ulong s)
3   {
4       for (ulong k=0; k<n; ++k)  s += a[k];
5       s %= nine;
6       return s;
7   }
```

where the variable `nine` has to be set to $R-1$ and `s` to zero.

The computation of $s_m = s_a \cdot s_b$ is done in

```
1   ulong
2   mult_sum_of_digits(const LIMB *a, ulong an,
3                      const LIMB *b, ulong bn,
4                      ulong nine)
5   {
6       ulong qsa = sum_of_digits(a, an, nine, 0);
7       ulong qsb = sum_of_digits(b, bn, nine, 0);
8       ulong qsm = (qsa*qsb) % nine;
9       return  qsm;
10  }
```

The checks in multiplication routine [FXT: mult/fxtmultiply.cc] can be outlined as:

```
1   fxt_multiply(const LIMB *a, ulong an,
2                const LIMB *b, ulong bn,
3                LIMB *c, ulong cn,
4                uint rx)
5   {
6       const ulong nine = rx-1;
7       ulong qsm=0, qsp=0;
8       qsm = mult_sum_of_digits(a, an, b, bn, nine);
9
10      // Multiply: c=a*b
11      // If carrying through c gives an additional (leading) digit,
12      // then set cy to that value,  else set cy=0.
13
14      qsp = sum_of_digits(g, n, nine, cy);
15      if ( qsm!=qsp )  { /* FAILED */ }
16  }
```

If we assume that a failed multiplication produces 'random' digits in c then the probability that a failed multiplication goes unnoticed equals $1/R$.

Omitting the sum-of-digits test is *not* an option: the situation that some number contains mainly 'nines' in the course of a high-precision calculation is very common. Thereby insufficient precision in the FFT-multiplication will almost certainly result in an error.

The simplicity of the sum-of-digits test that uses the modulus $R-1$ can be seen from the polynomial

identity

$$\sum_k a_k\, R^k \quad \equiv \quad \sum_k a_k \qquad \bmod R-1 \tag{27.5-1}$$

One can use other moduli, like

$$\sum_k a_k\, R^k \quad \equiv \quad \sum_k (-1)^k\, a_k \qquad \bmod R+1 \tag{27.5-2}$$

Moduli $R^n - 1$ for small n are especially convenient:

$$\sum_k a_k\, R^k \quad \equiv \quad \sum_{k\equiv 0 \bmod 2} a_k + R \sum_{k\equiv 1 \bmod 2} a_k \qquad \bmod R^2-1 \tag{27.5-3a}$$

$$\equiv \quad \sum_{k\equiv 0 \bmod 3} a_k + R \sum_{k\equiv 1 \bmod 3} a_k + R^2 \sum_{k\equiv 2 \bmod 3} a_k \qquad \bmod R^3-1 \tag{27.5-3b}$$

$$\equiv \quad \sum_{U=0}^{n-1} R^U \left(\sum_{k\equiv U \bmod n} a_k \right) \qquad \bmod R^n-1 \tag{27.5-3c}$$

One can keep the sums corresponding to the powers R^U in separate variables. Note that the probability of an unrecognized error is reduced to approximately $1/R^n$. The multiplication of the residues involves additional work proportional to n^2.

27.6 Binary exponentiation

The *binary exponentiation* (or *binary powering*) scheme is a method to compute the e-th power of a number a using about $\log(e)$ multiplications and squarings. The term 'number' can be replaced by about anything one can multiply. That includes integers, floating point numbers, polynomials, matrices, integer remainders modulo some modulus, polynomials modulo a polynomial and so on. In fact, the given algorithms work for any group: we do not need commutativity but $a^n \cdot a^m = a^{n+m}$ must hold (power-associativity).

27.6.1 Right-to-left powering

The algorithm uses the binary expansion of the exponent: let $e \geq 0$, write e the base 2 as $e = [e_j, e_{j-1}, \ldots, e_1, e_0]$, $e_i \in \{0, 1\}$. Then

$$a^e \quad = \quad a^{1\cdot e_0}\, a^{2\cdot e_1}\, a^{4\cdot e_2}\, a^{8\cdot e_3}\, \cdots\, a^{2^j\, e_j} \tag{27.6-1a}$$

$$= \quad 1\, (a^1)^{e_0}\, (a^2)^{e_1}\, (a^4)^{e_2}\, \cdots\, (a^{2^j})^{e_j} \tag{27.6-1b}$$

We initialize a variable t by one, generate the powers $s_i = a^{2^i}$ by successive squarings $s_i = s_{i-1}^2 = (a^{2^{i-1}})^2$ and multiply t by s_i if e_i equals one. The following C++ code computes the e-th power of the (double precision) number a:

```
1   double power_r2l(double a, ulong e)
2   {
3       double t = 1;
4       if ( e )
5       {
6           double s = a;
7           while ( 1 )
8           {
9               if ( e & 1 )  t *= s;
10              e /= 2;
11              if ( 0==e )  break;
12              s *= s;
```

```
13          }
14      }
15
16      return t;
17  }
```

A trivial optimization is to avoid the multiplication by 1 if the exponent is a power of two:

```
1   double power_r2l(double a, ulong e)
2   {
3       if ( 0==e )  return 1;
4
5       double s = a;
6       while ( 0==(e&1) )
7       {
8           s *= s;
9           e /= 2;
10      }
11
12      a = s;
13      while ( 0!=(e/=2) )
14      {
15          s *= s;
16          if ( e & 1 )  a *= s;
17      }
18      return  a;
19  }
```

The program [FXT: arith/power-r2l-demo.cc] shows the quantities that occur with the computation of $p = 2^{38}$:

```
arg 1: 2 == a  [number to exponentiate]  default=2
arg 2: 38 == e  [exponent]  default=38
 e = 1..11.
0                          2                        2
1                          4                        2
1                         16                       64
0                        256                       64
0                      65536                       64
1                 4294967296             274877906944
p=a**e = 274877906944
```

In the right-to-left powering scheme the exponent is scanned starting from the lowest bit.

27.6.2 Left-to-right powering

The left-to-right binary powering algorithm scans the exponent starting from the highest bits. We use the facts that $a^{2k} = (a^k)^2$ and $a^{2k+1} = (a^k)^2\, a$. Implementation is simple:

```
1   double power_l2r(double a, ulong e)
2   {
3       if ( 0==e )  return 1;
4       double s = a;
5       ulong b = highest_one(e);
6       while ( b>1 )
7       {
8           b >>= 1;
9           s *= s;
10          if ( e & b )  s *= a;
11      }
12      return s;
13  }
```

The program [FXT: arith/power-l2r-demo.cc] shows the quantities that occur with the computation of $p = 2^{38}$ when the left-to-right scan is used:

```
arg 1: 2 == a  [number to exponentiate]  default=2
arg 2: 38 == e  [exponent]  default=38
 e = 1..11.
1                          2
0                          4                        4
0                         16                       16
1                        256                      512
1                     262144                   524288
0             274877906944             274877906944
p=a**e = 274877906944
```

All multiplications apart from the squarings happen with the unchanged value of a. This is an advantage

if a is a small (integer) value so that the multiplications are cheap. As a slightly extreme example, if one computes $7^{7^7} \approx +0.3759823526783 \cdot 10^{695975}$ to full precision, then the left-to-right powering is about 3 times faster. If a is a full-precision number (and multiplication is done via FFTs) then the FFT of a only need to be computed once. Thereby all multiplications except for the first count a squarings. This technique is called *FFT caching*.

The given powering algorithms are good enough for most applications. There are schemes that improve further. For repeated power computations, especially for very large exponents the schemes based on the so-called *addition chains* lead to better algorithms, see [168]. The so-called 'flexible window powering method' is described and analyzed in [93]. A readable survey of exponentiation methods is given in [130].

Techniques for accelerating computations of factorials and binomial coefficients are described in [164].

27.6.3 Cost of binary exponentiation of full-precision numbers

e :	e (radix 2)	#S	#M	#F	#C
1:	`.....1`	0	0	0	
2:	`....1.`	1	0	2	
3:	`....11`	1	1	5	
4:	`...1..`	2	0	4	
5:	`...1.1`	2	1	7	
6:	`...11.`	2	1	7	
7:	`...111`	2	2	10	9
8:	`..1...`	3	0	6	
9:	`..1..1`	3	1	9	
10:	`..1.1.`	3	1	9	
11:	`..1.11`	3	2	12	11
12:	`..11..`	3	1	9	
13:	`..11.1`	3	2	12	11
14:	`..111.`	3	2	12	11
15:	`..1111`	3	3	15	13
16:	`.1....`	4	0	8	
17:	`.1...1`	4	1	11	
18:	`.1..1.`	4	1	11	
19:	`.1..11`	4	2	14	13
20:	`.1.1..`	4	1	11	

e :	e (radix 2)	#S	#M	#F	#C
21:	`.1.1.1`	4	2	14	13
22:	`.1.11.`	4	2	14	13
23:	`.1.111`	4	3	17	15
24:	`.11...`	4	1	11	
25:	`.11..1`	4	2	14	13
26:	`.11.1.`	4	2	14	13
27:	`.11.11`	4	3	17	15
28:	`.111..`	4	2	14	13
29:	`.111.1`	4	3	17	15
30:	`.1111.`	4	3	17	15
31:	`.11111`	4	4	20	17
32:	`1.....`	5	0	10	
33:	`1....1`	5	1	13	
34:	`1...1.`	5	1	13	
35:	`1...11`	5	2	16	15
36:	`1..1..`	5	1	13	
37:	`1..1.1`	5	2	16	15
38:	`1..11.`	5	2	16	15
39:	`1..111`	5	3	19	17
40:	`1.1...`	5	1	13	

Figure 27.6-A: Cost of binary powering of full-precision numbers for small exponents e in terms of squarings (#S), multiplications (#M) and FFTs (#F). If the left-to-right exponentiation algorithm with FFT caching needs less FFTs then the number is given under (#C).

With full-precision numbers the cost of binary powering is the same for both the left-to-right and the right-to-left algorithm. As an example, to raise x to the 26-th power, note that $e = 26 = 11010_2$ and we can write

$$x^{26} \;=\; x^{16} \cdot x^8 \cdot x^2 \;=\; (((x^2)^2)^2)^2 \cdot ((x^2)^2)^2 \cdot (x^2) \tag{27.6-2}$$

Here we need four squarings and two multiplications. In general one needs $\lfloor \log_2 e \rfloor$ squarings and $h(e)-1$ multiplications were $h(e)$ is the number of set bits in the binary expansion of e. Figure 27.6-A lists the cost of the exponentiation for small exponents e in terms of squarings and multiplications and, assuming a squaring costs two FFTs and multiplication three, in terms of FFTs. The table was created with the program [FXT: arith/power-costs-demo.cc].

Chapter 28

Root extraction

We describe methods to compute the inverse, square root, and higher roots of a given number. The computation of any of these costs just the equivalent of a few full-precision multiplications.

28.1 Division, square root and cube root

28.1.1 Inverse and division

The ordinary division algorithm is far too expensive for numbers of extreme precision. Instead one replaces the division $\frac{a}{d}$ by the multiplication of a with the inverse of d. The inverse of d is computed by finding a starting approximation $x_0 \approx \frac{1}{d}$ and then iterating

$$x_{k+1} \quad = \quad x_k + x_k\,(1 - d\,x_k) \tag{28.1-1}$$

until the desired precision is reached. The convergence is quadratic (second order), which means that the number of correct digits is doubled with each step: if $x_k = \frac{1}{d}(1+e)$ then $x_{k+1} = \frac{1}{d}\,\left(1-e^2\right)$.

Moreover, each step only requires computations with twice the number of digits that were correct at its beginning. Still better: the multiplication $x_k(\dots)$ needs only to be done with half of the current precision as it computes the correcting digits (which alter only the less significant half of the digits). Thus, at each step we have 1.5 multiplications of the current precision: one full precision multiplication for $d\,x_k$ plus a half precision multiplication for $x_k(\dots)$. The total work amounts to $1.5 + 1.5/2 + 1.5/4 + \dots = 1.5 \cdot \sum_{n=0}^{N} \frac{1}{2^n}$ which is less than 3 full precision multiplications. The cost of a multiplication is set to $\sim N$ for the estimates made here, this gives a realistic picture for large N. Together with the final multiplication a division costs as much as 4 multiplications.

The numerical example given in figure 28.1-A shows the first steps of the computation of an inverse starting from a two-digit initial approximation.

The achieved precision can be determined by the absolute value of $(1-d\,x_k)$. In hfloat, when the achieved precision is below a certain limit a third order correction is used to assure maximum precision at the last step:

$$x_{k+1} \quad = \quad x_k + x_k\,(1 - d\,x_k) + x_k\,(1 - d\,x_k)^2 \tag{28.1-2}$$

One should in general *not* use algebraically equivalent forms like $x_{k+1} = 2\,x_k - d\,x_k^2$ (for the second order iteration) because computationally there is a difference: cancellation can occur and the information on the achieved precision is not found easily.

$$
\begin{aligned}
d &:= 3.1415926 \\
x_0 &:= 0.31 \qquad \text{[initial 2-digit approximation for } 1/d\text{]} \\
\\
d \cdot x_0 &:= 3.141 \cdot 0.3100 = 0.9737 \\
y_0 &:= 1.000 - d \cdot x_0 = 0.02629 \\
x_0 \cdot y_0 &:= 0.3100 \cdot 0.02629 = 0.0081(49) \\
x_1 &:= x_0 + x_0 \cdot y_0 = 0.3100 + 0.0081 = 0.3181 \\
\\
d \cdot x_1 &:= 3.1415926 \cdot 0.31810000 = 0.9993406 \\
y_1 &:= 1.0000000 - d \cdot x_1 = 0.0006594 \\
x_1 \cdot y_1 &:= 0.31810000 \cdot 0.0006594 = 0.0002097(5500) \\
x_2 &:= x_1 + x_1 \cdot y_1 = 0.31810000 + 0.0002097 = 0.31830975 \\
\\
d \cdot x_2 &:= 3.1415926 \cdot 0.31830975 = 0.99999955 \\
y_2 &:= 1.0000000 - d \cdot x_2 = 0.00000014 \\
x_2 \cdot y_2 &:= 0.31830975 \cdot 0.00000014 = 0.000000044 \\
x_3 &:= x_2 + x_2 \cdot y_2 = 0.31830975 + 0.000000044 = 0.31830979399
\end{aligned}
$$

Figure 28.1-A: First steps of the computation of the inverse of π.

28.1.2 Inverse square root

Computation of inverse square roots can be done using a similar scheme: find a starting approximation $x_0 \approx \frac{1}{\sqrt{d}}$ then iterate

$$x_{k+1} \quad = \quad x_k + x_k \frac{(1 - d\,x_k^2)}{2} \tag{28.1-3}$$

until the desired precision is reached. Convergence is again second order: if $x_k = \frac{1}{\sqrt{d}}(1 + e)$ then

$$x_{k+1} \quad = \quad \frac{1}{\sqrt{d}} \left(1 - \frac{3}{2}e^2 - \frac{1}{2}e^3\right) \tag{28.1-4}$$

When the achieved precision is below a certain limit a third order correction should be applied:

$$x_{k+1} \quad = \quad x_k + x_k \frac{(1 - d\,x_k^2)}{2} + x_k \frac{3\,(1 - d\,x_k^2)^2}{8} \tag{28.1-5}$$

To compute the square root first compute $\frac{1}{\sqrt{d}}$ then a final multiply with d gives $\sqrt{d}$.

Similar considerations as above (with squaring considered as expensive as multiplication[1]) give an operation count of 4 multiplications for computing $\frac{1}{\sqrt{d}}$ and 5 for $\sqrt{d}$. Note that this algorithm is considerably better than iterating $x_{k+1} := \frac{1}{2}\,(x_k + \frac{d}{x_k})$ because no long divisions are involved.

A unified routine that implements the computation of the inverse a-th roots is given in [hfloat: src/hf/itiroot.cc]. The general form of the divisionless iteration for the a-th root of d is, up to third

[1] Indeed it costs about $\frac{2}{3}$ of a multiplication.

order:

$$x_{k+1} \quad = \quad x_k \left(1 + \frac{(1 - d\,x_k^a)}{a} + \frac{(1+a)\,(1 - d\,x_k^a)^2}{2\,a^2}\right) \tag{28.1-6}$$

The initial approximation is obtained using ordinary floating point numbers (type `double`) with special precautions to avoid overflow with exponents that cannot be represented with `double`s. Third order corrections are made whenever the achieved precision falls below a certain limit.

28.1.3 Cube root extraction

We use the relation $d^{1/3} = d\,(d^2)^{-1/3}$. That is, we compute the inverse third root of d^2 using the iteration

$$x_{k+1} \quad = \quad x_k + x_k \frac{(1 - d^2\,x_k^3)}{3} \tag{28.1-7}$$

and finally multiply with d. Convergence is second order: if $x_k = \frac{1}{\sqrt[3]{d}}(1+e)$ then

$$x_{k+1} \quad = \quad \frac{1}{\sqrt[3]{d}} \left(1 - 2e^2 - \frac{4}{3}e^3 - \frac{1}{3}e^4\right) \tag{28.1-8}$$

28.1.4 Improved iteration for the square root

Actually, the 'simple' version of the square root iteration ($x_{k+1} := \frac{1}{2}\,(x_k + \frac{d}{x_k})$) can be used for practical purposes when rewritten as a *coupled iteration* for both $\sqrt{d}$ and its inverse. Using for $\sqrt{d}$ the iteration

$$x_{k+1} \quad = \quad x_k - \frac{(x_k^2 - d)}{2\,x_k} \tag{28.1-9}$$

$$= \quad x_k - v_{k+1}\,\frac{(x_k^2 - d)}{2} \quad \text{where} \quad v \approx 1/x \tag{28.1-10}$$

and for the auxiliary $v \approx 1/\sqrt{d}$ the iteration

$$v_{k+1} = v_k + v_k\,(1 - x_k\,v_k) \tag{28.1-11}$$

where one starts with approximations

$$x_0 \quad \approx \quad \sqrt{d} \tag{28.1-12}$$

$$v_0 \quad \approx \quad 1/x_0 \tag{28.1-13}$$

and the v-iteration step precedes that for x. When carefully implemented this method turns out to be significantly more efficient than the computation via the inverse root. An implementation is given in [hfloat: src/hf/itsqrt.cc]. The idea is due to Schönhage.

28.1.5 A different view on the iterations

Let p be a prime and assume you know the inverse x_0 of a given number d modulo p. With (the iteration for the inverse, relation 28.1-1 on page 559) $\Phi(x) := x\,(1 + (1 - d\,x))$ the number $x_1 := \Phi(x_0)$ is the inverse of d modulo p^2. Modulo p^2 we know that $x_0\,d \equiv (1 + k\,p)$ so we can write $x_0 \equiv 1/d\,(1 + k\,p)$, thereby

$$\Phi(x_0) \quad = \quad \Phi\left(\frac{1}{d}\,(1 + k\,p)\right) \;=\; \frac{1}{d}\,\left(1 - k\,p^2\right) \;\equiv\; \frac{1}{d} \bmod p^2 \tag{28.1-14}$$

The very same computation (with $x_1 = 1/d\,(1 + j\,p^2)$) shows that for $x_2 := \Phi(x_1)$ one has $x_2 \equiv 1/d \bmod p^4$. Each application of Φ doubles the exponent of the modulus.

The equivalent scheme works for root extraction. An example for the inverse square root: with $p = 17$ and $x_0 = 3$ we have $x_0^2\,2 \equiv 1$ ($x_0 \equiv 1/\sqrt{d} \bmod p$ where $d = 2$). Now use the iteration $\Phi(x) := x\,(1+(1-d\,x)/2)$ to compute $x_1 = \Phi(x_0)$ to obtain $x_1 = -45/2 \equiv 122 \bmod p$ and observe that $x_1^2\,d \equiv 1 \bmod p^2$. Compute $x_2 = \Phi(x_1)$ to obtain $x_2 = -1815665 \equiv 21797 \bmod p^2$ and check that $x_2^2\,d \equiv 1 \bmod p^4$.

We will not go into the details (of the theory of p-adic numbers) but note that pari/gp can work with them:

```
? 1/sqrt(2+O(17^5))
  3 + 7*17 + 7*17^2 + 4*17^3 + 11*17^4 + O(17^5)
\\ Note that 21797 = 3 + 7*17 + 7*17^2 + 4*17^3 + 11*17^4
\\       and   122 = 3 + 7*17
```

Section 1.23 on page 59 describes the case $p = 2$. The computation of a square root modulo p^x, given a square root modulo p, is described in section 37.9 on page 773.

28.2 Root extraction for rationals

We give expression for the extraction of the a-th root of a rational quantity.

28.2.1 Extraction of the square root

A general formula for an k-th order ($k \geq 2$) iteration for $\sqrt{d}$ is

$$\Phi_k(x) \;=\; \sqrt{d}\,\frac{\left(x+\sqrt{d}\right)^k + \left(x-\sqrt{d}\right)^k}{\left(x+\sqrt{d}\right)^k - \left(x-\sqrt{d}\right)^k} = \sqrt{d}\,\frac{\left(p+q\sqrt{d}\right)^k + \left(p-q\sqrt{d}\right)^k}{\left(p+q\sqrt{d}\right)^k - \left(p-q\sqrt{d}\right)^k} \tag{28.2-1}$$

where $x = p/q$. All $\sqrt{d}$ vanish when expanded:

$$\Phi_2(x) \;=\; \frac{x^2+d}{2\,x} \;=\; \frac{p^2+d\,q^2}{2\,p\,q} \tag{28.2-2a}$$

$$\Phi_3(x) \;=\; x\,\frac{x^2+3\,d}{3\,x^2+d} \;=\; \frac{p}{q}\,\frac{p^2+3\,d\,q^2}{3\,p^2+d\,q^2} \tag{28.2-2b}$$

$$\Phi_4(x) \;=\; \frac{x^4+6\,d\,x^2+d^2}{4\,x^3+4\,d\,x} \;=\; \frac{p^4+6\,d\,p^2\,q^2+d^2\,q^4}{4\,p^3\,q+4\,d\,p\,q^3} \tag{28.2-2c}$$

$$\Phi_5(x) \;=\; x\,\frac{x^4+10\,d\,x^2+5\,d^2}{5\,x^4+10\,d\,x^2+d^2} \;=\; \frac{p}{q}\,\frac{p^4+10\,d\,p^2\,q^2+5\,d^2\,q^4}{5\,p^4+10\,d\,p^2\,q^2+d^2\,q^4} \tag{28.2-2d}$$

$$\Phi_k(x) \;=\; x\,\frac{\sum_{j=0}^{\lfloor k/2\rfloor}\binom{k}{2\,j}\,x^{k-2j}\,d^j}{\sum_{j=0}^{\lfloor k/2\rfloor}\binom{k}{2\,j+1}\,x^{k-2j-1}\,d^j} \tag{28.2-2e}$$

The denominators and numerators of Φ_k are terms of the second order recurrence

$$a_k \;=\; 2\,x\,a_{k-1} - \left(x^2-d\right)\,a_{k-2} \tag{28.2-3}$$

with initial terms $a_0 = 1$, $a_1 = x$ for the numerators, and $a_0 = 0$, $a_1 = 1$ for the denominators (that is, $\Phi_0 = 1/0$, $\Phi_1 = x/1$). An equivalent form of relation 28.2-1 is

$$\Phi_k(x) \;=\; \sqrt{d}\,\cot\left(k\,\operatorname{arccot}\frac{x}{\sqrt{d}}\right) \tag{28.2-4}$$

Setting $d = -1$ and $x = \cot(z)$ in relation 28.2-4 we obtain

$$\cot(k\,z) \quad = \quad \Phi_k(\cot(z)) \tag{28.2-5}$$

From this relation we obtain the following composition law:

$$\Phi_m(\Phi_n(x)) \quad = \quad \Phi_{m\,n}(x) \tag{28.2-6}$$

There is a nice expression for the error behavior of the k-th order iteration:

$$\Phi_k\left(\sqrt{d}\cdot\frac{1+e}{1-e}\right) \quad = \quad \sqrt{d}\cdot\frac{1+e^k}{1-e^k} \tag{28.2-7}$$

28.2.2 Extraction of the r-th root

A second order iteration for $\sqrt[r]{z}$ is given by

$$\Phi_2(x) \quad = \quad x + \frac{d-x^r}{r\,x^{r-1}} = \frac{(r-1)\,x^r + d}{r\,x^{r-1}} = \frac{1}{r}\left((r-1)\,x + \frac{d}{x^{r-1}}\right) \tag{28.2-8}$$

A third order iteration for $\sqrt[r]{d}$ is

$$\Phi_3(x) \quad = \quad x\cdot\frac{\alpha\,x^r + \beta\,d}{\beta\,x^r + \alpha\,d} \;=\; \frac{p}{q}\cdot\frac{\alpha\,p^r + \beta\,q^r d}{\beta\,p^r + \alpha\,q^r d} \tag{28.2-9}$$

where $x = p/q$, $\alpha = r-1$ and $\beta = r+1$. An alternative form is

$$\Phi_3(x) \quad = \quad \frac{d}{x}\cdot\frac{\beta\,x^r + \alpha\,d}{\alpha\,x^r + \beta\,d} \tag{28.2-10}$$

where again $\alpha = r-1$ and $\beta = r+1$.

28.2.3 More iterations via Padé approximants

The iterations can also be obtained using Padé-approximants. Let $P_{[i,j]}(z)$ be the Padé-expansion of $\sqrt{z}$ around $z = 1$ of order $[i,j]$. An iteration of order $i+j+1$ is given by $x\,P_{[i,j]}(\frac{d}{x^2})$. Different combinations of i and j result in alternative iterations:

$$[i,j] \quad \mapsto \quad x\,P_{[i,j]}\left(\frac{d}{x^2}\right) \tag{28.2-11a}$$

$$[1,0] \quad \mapsto \quad \frac{x^2+d}{2x} \tag{28.2-11b}$$

$$[0,1] \quad \mapsto \quad \frac{2x^3}{3x^2-d} \tag{28.2-11c}$$

$$[1,1] \quad \mapsto \quad x\,\frac{x^2+3d}{3x^2+d} \tag{28.2-11d}$$

$$[2,0] \quad \mapsto \quad \frac{3x^4+6dx^2-3d^2}{8x^3} \tag{28.2-11e}$$

$$[0,2] \quad \mapsto \quad \frac{8x^5}{15x^4-10dx^2+3d^2} \tag{28.2-11f}$$

Still other forms are obtained by using $\frac{d}{x}\,P_{[i,j]}(\frac{x^2}{d})$:

$$[i,j] \;\mapsto\; \frac{d}{x}\,P_{[i,j]}\left(\frac{x^2}{d}\right) \tag{28.2-12a}$$

$$[1,0] \;\mapsto\; \frac{x^2+d}{2x} \tag{28.2-12b}$$

$$[0,1] \;\mapsto\; \frac{2d^2}{3dx-x^3} \tag{28.2-12c}$$

$$[1,1] \;\mapsto\; \frac{d\,(d+3x^3)}{x\,(3d+x^2)} \tag{28.2-12d}$$

$$[2,0] \;\mapsto\; \frac{-x^4+6dx^2+3d^2}{8xd} \tag{28.2-12e}$$

$$[0,2] \;\mapsto\; \frac{8d^3}{3x^4-10dx^2+15d^2} \tag{28.2-12f}$$

Pade approximations for the r-th root can be given in terms of hypergeometric functions. Rewrite relation 35.2-30 on page 689 as

$$F\left(\begin{matrix}u,\,v\\ u+v+1/r\end{matrix}\,\middle|\,z\right)/F\left(\begin{matrix}u+1/r,\,v+1/r\\ u+v+1/r\end{matrix}\,\middle|\,z\right) \;=\; (1-z)^{1/r} \tag{28.2-13}$$

The expression on the left gives the approximant $[i,j]$ if we set $u=-i$ and $v=-j-1/r$ (so both series terminate). An iteration of order $i+j+1$ is obtained as $\Phi_{[i,j]} = x\,P_{[i,j]}\left(\frac{d}{x^r}\right)$. We compute the third order iteration for the fourth root (relation 28.2-9 with $r=4$):

```
? r=4;  \\ r-th root
? i=1;  \\ degree of denominator
? j=1;  \\ degree of numerator
? u=-i;v=-j-1/r;  \\ setup parameters so that series terminate
? N=hypergeom([u,v],[u+v+1/r],x,i)
  -5/8*x + 1
? D=hypergeom([u+1/r,v+1/r],[u+v+1/r],x,j)
  -3/8*x + 1
? t=N/D  \\ Pade approximant [i,j]
  (-5*x + 8)/(-3*x + 8)

\\ check t == (1-x)^(1/r) + order(i+j+1):
? n=i+j+2;
? t-hypergeom([-1/r],[],x,n)+O(x^(n))
  5/256*x^3 + O(x^4)

? t=subst(t,x,1-x)  \\ Pade approximant in x
  (5*x + 3)/(3*x + 5)
? it=x*subst(t,x,d/x^r)  \\ iteration for f(x)=x^r-d
  (3*x^5 + 5*d*x)/(5*x^4 + 3*d)
\\ ==   x * (3*x^4 + 5*d)/(5*x^4 + 3*d)
```

Now we check the order of the iteration, we set $d=1$ and compute $f\left(\Phi\left(d^{1/r}(1+e)\right)\right) = f\left(\Phi(1+e)\right)$:

```
? f(x)=x^r-1;  \\ ==x^r-d  for d==1
? it=subst(it,d,1);  \\ d==1
? er=subst(it,x,(1+e));  \\ Phi( d^(1/r))*(1+e) ) = Phi(1+e)
? taylor(f(er),e)  \\ f( Phi(1+e) ) =?= O(e^(i+j+1))
  5*e^3 - 15/2*e^4 + [...]  \\ OK
```

An alternative form of iterations can be obtained as $\Phi_{[i,j]} = \frac{d}{x}\,P_{[i,j]}\left(\frac{x^r}{d}\right)$. The $[1,1]$-approximant leads to the third order iteration given as relation 28.2-10.

In section 29.6.1 on page 589 Padé approximants are used to obtain iterations for arbitrary functions f.

28.3 Divisionless iterations for the inverse a-th root

There is a nice general formula that gives iterations with arbitrary order of convergence for $1/\sqrt[a]{d} = d^{-1/a}$ that involve no long division.

One uses the identity

$$d^{-1/a} = x\,(1-(1-x^a\,d))^{-1/a} \tag{28.3-1}$$
$$= x\,(1-y)^{-1/a} \quad \text{where} \quad y := (1-x^a\,d) \tag{28.3-2}$$

Taylor expansion gives

$$d^{-1/a} = x\sum_{k=0}^{\infty}(1/a)^{\overline{k}}\,y^k \tag{28.3-3}$$

where $z^{\overline{k}} := z\,(z+1)\,(z+2)\,\ldots\,(z+k-1)$ (and $z^{\overline{0}} := 1$, $z^{\overline{k}}$ is called the rising factorial power). Written out:

$$d^{-1/a} = x\,\frac{1}{\sqrt[a]{1-y}} = x\left(1+\frac{y}{a}+\frac{(1+a)\,y^2}{2\,a^2}+\frac{(1+a)(1+2a)\,y^3}{6\,a^3}+\right. \tag{28.3-4}$$
$$\left.+\frac{(1+a)(1+2a)(1+3a)\,y^4}{24\,a^4}+\cdots+\frac{\prod_{k=1}^{n-1}(1+k\,a)}{n!\,a^n}\,y^n+\ldots\right)$$

A n-th order iteration for $d^{-1/a}$ is obtained by truncating the above series after the $(n-1)$-th term:

$$\Phi_n(x) := x\sum_{k=0}^{n-1}(1/a)^{\overline{k}}\,y^k \tag{28.3-5}$$
$$x_{k+1} = \Phi_n(x_k) \tag{28.3-6}$$

Convergence is n-th order:

$$\Phi_n(d^{-1/a}(1+e)) = d^{-1/a}(1+O(e^n)) \tag{28.3-7}$$

For example, the second order iteration is

$$\Phi_2(x) := x + x\frac{(1-d\,x^a)}{a} \tag{28.3-8}$$

Convergence is indeed quadratic: if $x = \frac{1}{\sqrt[a]{d}}\,(1+e)$ then

$$\Phi_2(x) = \frac{1}{\sqrt[a]{d}}\left((1+e)\left[(1+e)^a-(a+1)\right]\right) \tag{28.3-9}$$
$$= \frac{1}{\sqrt[a]{d}}\left(1-\frac{a+1}{2}e^2+O(e^3)\right) \tag{28.3-10}$$

28.3.1 Iterations for the inverse

Set $a=1$, $y=1-d\,x$ to compute the inverse of d.

$$\frac{1}{d} = x\,\frac{1}{1-y} \tag{28.3-11a}$$
$$\Phi_k(x) = x\,\left(1+y+y^2+y^3+y^4+\cdots+y^{k-1}\right) \tag{28.3-11b}$$

$\Phi_2(x) = x\,(1+y)$ is the second order iteration 28.1-1 on page 559.

Composition is particularly simple with the iterations for the inverse:

$$\Phi_{n\,m} \;=\; \Phi_n(\Phi_m) \tag{28.3-12}$$

There are simple closed forms for this iteration:

$$\Phi_k \;=\; \frac{1-y^k}{d} = x\,\frac{1-y^k}{1-y} \tag{28.3-13a}$$

$$\Phi_\infty \;=\; 1 + x + x^2 + x^3 + x^4 + \ldots \tag{28.3-13b}$$

$$\;=\; x\,(1+y)\,(1+y^2)\,(1+y^4)\,(1+y^8)\,\ldots \tag{28.3-13c}$$

$$\;=\; x\,(1+y+y^2)\,(1+y^3+y^6)\,(1+y^9+y^{18})\,\ldots \tag{28.3-13d}$$

The expression for the convergence of the k-th order iteration is

$$\Phi_k\left(\frac{1}{d}\,(1+e)\right) \;=\; \frac{1}{d}\,\left(1-(-e)^k\right) \tag{28.3-14}$$

The iteration converges if one has $|e| < 1$ for the start value $x_0 = \frac{1}{d}\,(1+e)$. That is, the basin of attraction is the open disc of radius $r = 1/d$ around the point $1/d$, independent of the order k. For other iterations, the basin of attraction usually has a fractal boundary and further depends on the order.

28.3.2 Iterations for the inverse square root

Set $a = 2$, $y = 1 - d\,x^2$ to compute the inverse square root of d.

$$\frac{1}{\sqrt{d}} \;=\; x\,\frac{1}{\sqrt{1-y}} \tag{28.3-15a}$$

$$\;=\; x\left(1 + \frac{y}{2} + \frac{3\,y^2}{8} + \frac{5\,y^3}{16} + \frac{35\,y^4}{128} + \cdots + \frac{\binom{2k}{k}\,y^k}{4^k} + \ldots\right) \tag{28.3-15b}$$

$$\Phi_{k+1}(x) \;=\; x\left(1 + \frac{y}{2} + \frac{3\,y^2}{8} + \ldots + \frac{\binom{2k}{k}\,y^k}{4^k}\right) \tag{28.3-15c}$$

$\Phi_2(x) = x\,(1+y/2)$ is the second order iteration 28.1-3 on page 560.

28.3.3 Computation of the a-th root

The following (second order) iteration computes $\sqrt[a]{d}$ directly:

$$\Phi(x) \;=\; x + \frac{1}{a}\left(x - \frac{x^{a+1}}{d}\right) \tag{28.3-16}$$

Figure 28.3-A shows the quantities occurring in the iterative computation of $\sqrt[5]{3}$. The iteration is involves no long division for small (rational) d.

To compute the a-root of a full-precision number d one can use the iteration for the inverse root and invert afterwards. Another possibility is to compute the inverse a-th root of d^{a-1} and multiply with d afterwards:

$$\left[(d)^{a-1}\right]^{-1/a}\,d \;=\; d^{(1-a)/a}\,d \;=\; d^{1/a} \tag{28.3-17}$$

If a is small the cost is lower than with the final iteration for the inverse (which costs about 3 Multiplications or 9 FFTs). The powering-method is not more expensive than inversion if the rightmost column in figure 27.6-A on page 557 for $e = a - 1$ is less or equal 6. In case the iteration for the inverse involves a loss of precision the method might be preferred also if its cost is higher.

```
? default(realprecision,55);
? n=5;d=3;
? f=x^n-d
   x^5 - 3
? phi(x)=(x+(x/n-x^(n+1)/(n*d)))
? phi(x)
   -1/15*x^6 + 6/5*x
? y=real(polroots(f)[1])
   1.24573093961551732596668033664030508093930999306877981 1
? y*=(1.01); \\ <--= initial approximation within 1%
? for(k=0,7,t=phi(y);print(k,":  ",y);y=t;);
0:  1.25818824901167249922634714000670813174870309299946760 9
1:  1.24535219988820916129228150423636135252138734392268204 9
2:  1.24573059431066513233812676008483214085083362188066730 3
3:  1.24573093961523018034055334316279342578350028013835627 2
4:  1.24573093961551732596668013807589285894062440312487496 2
5:  1.24573093961551732596668033664030508093930999306868486 0
6:  1.24573093961551732596668033664030508093930999306877981 1
7:  1.24573093961551732596668033664030508093930999306877981 1
? y^n
   3.000000000000000000000000000000000000000000000000000000
```

Figure 28.3-A: Quantities occurring in the iterative computation of $\sqrt[5]{3}$.

28.3.4 Error expressions for inverse square root iterations *

An expression for the error behavior of the n-th order iteration similar to relation 28.2-7 on page 563 is

$$F_n := \Phi_n(d^{-1/2}\frac{1+e}{1-e})/d^{-1/2} \tag{28.3-18a}$$

$$= \frac{\sum_{k=0}^{n}\binom{2n-1}{k}(-e)^k - \sum_{k=n+1}^{2n+1}\binom{2n-1}{k}(-e)^k}{(1-e)^{2n-1}} \tag{28.3-18b}$$

Now define $c := \frac{F_n-1}{F_n+1}$, then

$$F_n = \frac{1+c}{1-c} \quad \text{where} \quad c = e^n \frac{\sum_{k=0}^{n}\binom{2n-1}{n-k}(-e)^k}{\sum_{k=0}^{n}\binom{2n-1}{k}(-e)^k} \tag{28.3-18c}$$

For example, with $n = 2$ we obtain

$$F_2 := \Phi_2(d^{-1/2}\frac{1+e}{1-e}) \,/\, d^{-1/2} \tag{28.3-19a}$$

$$= \frac{1-3\,e-3\,e^2+e^3}{1-3\,e+3\,e^2-e^3} \tag{28.3-19b}$$

$$= 1-6\,e^2-16\,e^3-30\,e^4-48\,e^5-70\,e^6-\ldots \tag{28.3-19c}$$

$$= -1+\frac{6}{(e-1)^2}+\frac{4}{(e-1)^3} \tag{28.3-19d}$$

$$F_2 = \frac{1+c}{1-c} \quad \text{where} \quad c = e^2\,\frac{e-3}{1-3\,e} \tag{28.3-19e}$$

The coefficients of the Taylor expansion of F_n in e are always integers.

For $n = 4$:

$$F_4 := \Phi_4(d^{-1/2}\frac{1+e}{1-e}) \,/\, d^{-1/2} \tag{28.3-20a}$$

$$= \frac{1 - 7\,e + 21\,e^2 - 35\,e^3 - 35\,e^4 + 21\,e^5 - 7\,e^6 + e^7}{1 - 7\,e + 21\,e^2 - 35\,e^3 + 35\,e^4 - 21\,e^5 + 7\,e^6 - e^7} \tag{28.3-20b}$$

$$= 1 - 70\,e^4 - 448\,e^5 - 1680\,e^6 - 4800\,e^7 - 11550\,e^8 - \ldots \tag{28.3-20c}$$

$$= -1 + \frac{70}{(e-1)^4} + \frac{168}{(e-1)^5} + \frac{140}{(e-1)^6} + \frac{40}{(e-1)^7} \tag{28.3-20d}$$

$$F_4 = \frac{1+c}{1-c} \quad \text{where} \quad c = e^4\,\frac{e^3 - 7\,e^2 + 21\,e - 35}{1 - 7\,e + 21\,e^2 - 35\,e^3} \tag{28.3-20e}$$

Two curious formulas related to the error behavior of Φ_2 are

$$\Phi_2\left(\frac{1}{\sqrt{d}}\left[e + \frac{1}{e}\right]\right) = \frac{1}{\sqrt{d}}\left[-\frac{1}{2}\cdot\left(e^3 + \frac{1}{e^3}\right)\right] \tag{28.3-21}$$

$$\Phi_2\left(\frac{1}{\sqrt{d}}\left[e - \frac{2}{3}\,\frac{1}{e}\right]\right) = \frac{1}{\sqrt{d}}\left[+\frac{1}{2}\cdot\left(e^3 - \frac{2}{3}\,\frac{1}{e^3}\right)\right] \tag{28.3-22}$$

28.4 Initial approximations for iterations

With the iterative schemes one alway needs an initial approximation for the root that is to be computed. Assume we want to compute $f(d)$, for example, $f(d) = \sqrt{d}$ or $f(d) = \exp(d)$. One can convert the high precision number d to a machine floating point number and use the FPU to compute an initial approximation. However, when d cannot be represented with a machine float, the method fails. The method will also fail if the result causes an overflow, which is likely to happen with $f(d) = \exp(d)$.

28.4.1 Inverse roots

With $f(d) = d^{1/a}$ one can use the following technique. Write d in the form

$$d = M \cdot R^X \tag{28.4-1}$$

were M is the mantissa, R the radix and X the exponent. We have $0 \le M < 1$ and $X \in \mathbb{Z}$. Now use

$$d^{1/a} = M^{1/a} \cdot R^{X/a} = M^{1/a} \cdot R^{Y/a} \cdot R^Z \tag{28.4-2}$$

where $Z = \lfloor X/a \rfloor$ and $Y = X - a \cdot Z$ (so $X = a \cdot Z + Y$).

The algorithm computes the three quantities in relation 28.4-2 separately and finally computes the product as result. A C++ implementation is given in [hfloat: src/hf/itiroot.cc]:

```
1   void
2   approx_invpow(const hfloat &d, hfloat &c, long a)
3   {
4       double dd;
5       dt_mantissa_to_double(*(d.data()), dd);
6       dd = pow(dd, 1.0/(double)a);  // M^(1/a)
7
8       long  Z = d.exp() / a;    // Z = X / a
9       long  Y = d.exp() - a*Z;  // Y = X % a
10
11      double tt = pow((double)d.radix(),(double)Y/a);  // R^(Y/a)
12      dd *= tt;          // M^(1/a) * R^(Y/a)
13
14      d2hfloat(dd, c);  // c =  M^(1/a) * R^(Y/a)
15      c.exp( c.exp()+Z );   // c *= R^(Z)
16  }
```

One can also subtract $a \cdot Z$ from the exponent before the iteration and add Z to the exponent afterwards: $\left(R^{X-a\,Z}\right)^{1/a} = R^{Y/a} = R^{X/a}/R^{Z}$. In that case the initial approximation can be computed via the straight forward approach.

28.4.2 Exponential function

With $f(d) = \exp(d)$ write

$$\exp(d) \quad = \quad M \cdot R^{X}$$

Then use $X = \lfloor d/\log(R) \rfloor$ and $M = \exp(d - X \cdot \log R)$. Note that d must fit into a machine float which is not a real restriction as the $\exp(d)$ will not fit into a `hfloat` type already with smaller values as the exponent of the result would overflow.

A C++ implementation is given in [hfloat: src/tz/itexp.cc]:

```
1   void
2   approx_exp(const hfloat &d, hfloat &c)
3   {
4       double dd;
5       hfloat2d(d,dd);
6
7       double lr = log( hfloat::radix() );
8       double X = floor( dd/lr );
9       double M = exp( dd-X*lr );
10
11      d2hfloat(M,c);
12      c.exp( c.exp()+(long)X );
13  }
```

The iteration for computation of the exponential function is given in section 31.2 on page 623.

28.5 Some applications of the matrix square root

We give applications of the iteration for the (inverse) square root to compute re-orthogonalized matrices, the polar decomposition, the sign decomposition, and the pseudo-inverse of a matrix.

28.5.1 Re-orthogonalization

A task from graphics applications: a rotation matrix A that deviates from being orthogonal[2] shall be transformed to the closest orthogonal matrix E. One has (see [228]):

$$E \quad = \quad A\,(A^T A)^{-\frac{1}{2}} \tag{28.5-1}$$

With the division-free iteration for the inverse square root

$$\Phi(x) \quad = \quad x\left(1 + \frac{1}{2}\left(1 - dx^2\right) + \frac{3}{8}\left(1 - dx^2\right)^2 + \frac{5}{16}\left(1 - dx^2\right)^3 + \ldots\right) \tag{28.5-2}$$

the given task is pretty easy: as $A^T A$ is close to unity (the identity matrix) we can use the (second order) iteration with $d = A^T A$ and $x = 1$

$$(A^T A)^{-\frac{1}{2}} \quad \approx \quad \left(1 + \frac{1 - A^T A}{2}\right) \tag{28.5-3}$$

[2] Typically due to cumulative errors resulting from many multiplications with rotation matrices.

and multiply by A to get a 'closer-to-orthogonal' matrix A_+:

$$A_+ \;=\; A\left(1+\frac{1-A^TA}{2}\right) \;\approx E \tag{28.5-4}$$

The step can be repeated with A_+ (or higher orders can be used) if necessary. Note the identical equation would be obtained when trying to compute the inverse square root of 1:

$$x_+ \;=\; x\left(1+\frac{1-x^2}{2}\right) \;\to 1 \tag{28.5-5}$$

It is instructive to write things down in the *singular value decomposition* (SVD) representation

$$A \;=\; U\,\Omega\,V^T \tag{28.5-6}$$

where U and V are orthogonal and Ω is a diagonal matrix with non-negative entries, see [280]. We note that the SVD is *not* unique, using the 1×1 matrix $[-2]=[-1]\,[2]\,[1]=[1]\,[2]\,[-1]$. The SVD is a decomposition of the action of the matrix as: rotation – element-wise stretching – rotation. Now

$$A^TA \;=\; \left(V\Omega U^T\right)\,\left(U\Omega V^T\right)=V\Omega^2V^T \tag{28.5-7}$$

Thereby (using the fact that $(V\Omega V^T)^n=V\Omega^nV^T$)

$$(A^TA)^{-\frac{1}{2}} \;=\; \left(\left(V\Omega U^T\right)\left(U\Omega V^T\right)\right)^{-\frac{1}{2}}=\left(V\Omega^2V^T\right)^{-\frac{1}{2}}=V\Omega^{-1}V^T \tag{28.5-8}$$

and we have

$$A\,(A^TA)^{-\frac{1}{2}} \;=\; \left(U\Omega V^T\right)\,\left(V\Omega^{-1}V^T\right)=U\,V^T \tag{28.5-9}$$

that is, the 'stretching part' was removed.

A numerical example: Let

$$A \;=\; \begin{bmatrix}+1.0000000 & +1.0000000 & +0.7500000\\ -0.5000000 & +1.5000000 & +1.0000000\\ +0.7500000 & +0.5000000 & -1.0000000\end{bmatrix} \tag{28.5-10}$$

then

$$E \;=\; \begin{bmatrix}+0.803114165 & +0.291073143 & +0.519888513\\ -0.486897253 & +0.823533541 & +0.291073143\\ +0.343422053 & +0.486897252 & -0.803114166\end{bmatrix} \tag{28.5-11}$$

and $E\,E^T=1$.

28.5.2 Polar decomposition

The so-called *polar decomposition* of a matrix A is a representation of the form

$$A \;=\; E\,R \tag{28.5-12}$$

where the matrix E is orthogonal and $R=R^T$. It is an analogue to the representation of a complex number $z\in\mathbb{C}$ as $z=e^{i\,\phi}\,r$ (identify $R\sim r$ and $E\sim e^{i\,\phi}$).

The polar decomposition can be defined by

$$A \;=\; E\,R \;:=\; \left(A(A^TA)^{-1/2}\right)\,\left((A^TA)^{1/2}\right) \tag{28.5-13}$$

where R is the (unique, positive semidefinite) square root $R = (A^T A)^{1/2}$ and $E = A(A^T A)^{-1/2}$.

The matrix E is computed as before:

$$E = A \cdot \left(1 + \frac{1 - A^T A}{2}\right) \cdot \left(1 + \frac{1 - A_+^T A_+}{2}\right) \cdot \ldots \tag{28.5-14}$$

The matrix R equals $E^{-1}A = E^T A$, that is

$$A = E\,R = E\,\left(E^T A\right) \tag{28.5-15}$$
$$= U\,V^T\,\left(V\,\Omega\,V^T\right) \tag{28.5-16}$$

Compute the polar decomposition as

$$E_0 = A \tag{28.5-17a}$$
$$Y_k = \left(1 + \frac{1 - E_k^T\,E_k}{2}\right) \tag{28.5-17b}$$
$$E_{k+1} = E_k\,Y_k \quad \to E \tag{28.5-17c}$$
$$R_{k+1} = E_{k+1}^T A \quad \to R \tag{28.5-17d}$$
$$E_{k+1}\,R_{k+1} \to A \tag{28.5-17e}$$

Higher orders can be added in the computation of Y_k. If you prefer $z = r\,e^{i\,\phi}$ over $e^{i\,\phi}\,r$ then iterate as above but set $R' = A\,E^T$ so that

$$A = R'\,E\,=\,\left(AE^T\right)\,E \tag{28.5-18}$$
$$= \left(U\,\Omega\,U^T\right)\,U\,V^T \tag{28.5-19}$$

Numerical example: Let

$$A = \begin{bmatrix} +1.00000 & +1.00000 & +0.75000 \\ -0.50000 & +1.50000 & +1.00000 \\ +0.75000 & +0.50000 & -1.00000 \end{bmatrix} \tag{28.5-20}$$

then

$$A = E\,R \tag{28.5-21a}$$
$$= \begin{bmatrix} +0.80311 & +0.29107 & +0.51988 \\ -0.48689 & +0.82353 & +0.29107 \\ +0.34342 & +0.48689 & -0.80311 \end{bmatrix} \begin{bmatrix} +1.30412 & +0.24447 & -0.22798 \\ +0.24447 & +1.76982 & +0.55494 \\ -0.22798 & +0.55494 & +1.48410 \end{bmatrix} \tag{28.5-21b}$$
$$A = R'\,E \tag{28.5-21c}$$
$$= \begin{bmatrix} +1.48410 & +0.55494 & +0.22798 \\ +0.55494 & +1.76982 & -0.24447 \\ +0.22798 & -0.24447 & +1.30412 \end{bmatrix} \begin{bmatrix} +0.80311 & +0.29107 & +0.51988 \\ -0.48689 & +0.82353 & +0.29107 \\ +0.34342 & +0.48689 & -0.80311 \end{bmatrix} \tag{28.5-21d}$$

28.5.3 Sign decomposition

The *sign decomposition* can be defined as

$$A = S\,N \quad = \left(A(A^2)^{-1/2}\right)\,\left((A^2)^{1/2}\right) \tag{28.5-22}$$

where $N = (A^2)^{-1/2}$ and $S = A\,(A^2)^{-1/2}$. The square root has to be chosen such that all its eigenvalues have positive real parts. The sign decomposition is undefined if A has eigenvalues on the imaginary axis. The matrix S is its own inverse (its eigenvalues are ± 1). The matrices A, S and N commute pairwise: $SN = NS$, $AN = NA$ and $AS = SA$.

Use

$$\begin{aligned} S_0 &= A && (28.5\text{-}23a)\\ Y_k &= \left(1+\frac{1-S_k^2}{2}\right) && (28.5\text{-}23b)\\ S_{k+1} &= S_k\, Y_k \quad \to S && (28.5\text{-}23c)\\ N_{k+1} &= S_{k+1}\, A \quad \to N && (28.5\text{-}23d) \end{aligned}$$

Numerical example: Let

$$A = \begin{bmatrix} +1.00000 & +1.00000 & +0.75000 \\ -0.50000 & +1.50000 & +1.00000 \\ +0.75000 & +0.50000 & -1.00000 \end{bmatrix} \tag{28.5-24}$$

then

$$\begin{aligned} A &= S\,N && (28.5\text{-}25a)\\ &= \begin{bmatrix} +0.90071 & -0.01706 & +0.29453 \\ -0.24065 & +0.95862 & +0.71389 \\ +0.62679 & +0.10775 & -0.85933 \end{bmatrix} \begin{bmatrix} +1.13014 & +1.02237 & +0.36392 \\ -0.18454 & +1.55423 & +0.06423 \\ -0.07158 & +0.35875 & +1.43718 \end{bmatrix} && (28.5\text{-}25b) \end{aligned}$$

where $S\,S = 1$. See [145] and also [146].

28.5.4 Pseudo-inverse

While we are at it: define a matrix A^+ as

$$A^+ := (A\,A^T)^{-1}\,A^T = \left(V\,\Omega^{-2}\,V^T\right)\,\left(V\,\Omega\,U^T\right) = V\,\Omega^{-1}\,U^T \tag{28.5-26}$$

This looks suspiciously like the inverse of A. In fact, this is the *pseudo-inverse* of A:

$$A^+\,A = \left(V\,\Omega^{-1}\,U^T\right)\,\left(U\,\Omega\,V^T\right) = 1 \quad \text{but wait} \tag{28.5-27}$$

A^+ has the nice property to exist even if A^{-1} does not. If A^{-1} exists, it is identical to A^+. If not, $A^+\,A \neq 1$ but A^+ will give the best possible (in a least-square sense) solution $x^+ = A^+ b$ of the equation $A\,x = b$ (see [95, p.770]). To find $(A\,A^T)^{-1}$ use the iteration for the inverse:

$$\Phi(x) = x\,\left(1+(1-dx)+(1-dx)^2+\dots\right) \tag{28.5-28}$$

with $d = A\,A^T$ and the start value $x_0 = 2 - n\,(A\,A^T)/\,||A\,A^T||^2$ where n is the dimension of A.

A pari/gp implementation of the pseudo-inverse using the SVD:

```
1   matpseudoinv(A)=
2   \\ Return pseudo-inverse of A
3   {
4       local(t, x, U, d, V);
5       t = matSVD(A);
6       U = t[1];  d = t[2];  V = t[3];
7       for (k=1, matsize(d)[1],
8           x=d[k,k];  if (x>1e-15, d[k,k]=1/x, d[k,k]=0);
9       );
10      return( V*d*U~ );
11  }
```

Where the SVD is computed with the help of a routine (`qfjacobi()`) that returns the eigenvectors of a real symmetric matrix:

```
? A
  [+1.00 +1.00 +0.75 +2.00]
  [-0.50 +1.50 +1.00 +3.00]
  [+0.75 +0.50 -1.00 -3.00]

? t=matSVD(A); U=t[1]; d=t[2]; V=t[3];
? U
  [+0.644401153492 +0.438818890 +0.6262468643]
  [-0.695372132379 +0.676976586 +0.2411644651]
  [-0.318126941467 -0.590881276 +0.7413869203]

? d
  [+0.95641003 0 0]
  [0 +5.09161169 0]
  [0 0 +1.74234618]

? V
  [+0.787833655771 -0.067332385426 +0.60935354299]
  [-0.583139548860 +0.227598489665 +0.77980313700]
  [+0.110889336609 +0.313647647860 -0.01752654388]
  [+0.164225309908 +0.919396775264 -0.14243646797]

? Ax=matpseudoinv(A)
  [+0.744034618880 -0.497415792829 +0.005046325813]
  [-0.093004327360 +0.562176974103 +0.499369209273]
  [+0.095446097914 -0.041347314730 -0.080741213017]
  [+0.138692567521 -0.016875347613 -0.221929812661]

? A*Ax
  [+1.0000000000000 +3.78653234 E-29 -4.41762106 E-29]
  [+2.52435489 E-29 +1.0000000000000 -2.52435489 E-29]
  [-2.52435489 E-29 -2.52435489 E-29 +1.0000000000000]

? Ax*A
  [+0.9965272596551 +0.0004340925431 +0.0555638455173 -0.0193171181681]
  [+0.0004340925431 +0.9999457384321 -0.0069454806896 +0.0024146397710]
  [+0.0555638455173 -0.0069454806896 +0.1109784717229 +0.3090738906900]
  [-0.0193171181681 +0.0024146397710 +0.3090738906900 +0.8925485301897]
```

Figure 28.5-A: Numerical example for the pseudo-inverse computed by the SVD. We use a 3×4 matrix which is definitely not invertible. A working precision of 25 decimal digits was used, so $A\,A^{+} = 1$ to within that precision. On the other hand, $A^{+}\,A$ is not close to the unit matrix.

```
1   matSVDcore(A)=
2   \\ Singular value decomposition:
3   \\ Return [U, d, V] so that U*d*V~==A
4   \\ d is a diagonal matrix
5   \\ U, V are orthogonal
6   {
7       local(U, d, V);  \\ returned quantities
8       local(t, R, d1);
9       R = conj(A~)*A;  \\ R==V*d^2*V~
10      t = qfjacobi( R );  \\ fails with eigenvalues==zero
11      V = t[2];
12      d = real(sqrt(t[1]));
13      d1 = d;
14      for (k=1, length(d1), t=d1[k]; if (abs(t)>1e-16, t=1/t, t=0); d1[k]=t );
15      d1 = matdiagonal(d1):
16      d = matdiagonal(d);
17      U = (A*V*d1);
18      return( [U, d, V] );
19  }
```

The core routine is always called with a matrix A whose number of rows is greater or equal to its number of rows. Thereby we make sure that when the main routine with argument A computes U, d, V so that $A = U\,d\,V^{T}$ then with argument A^{T} we obtain X, d, Y where $A^{T} = X\,d\,Y^{T}$, $X = V$, and $Y = U$.

```
1   matSVD(A)=
2   {
3       local(tq, t, U, d, V);
4       t = matsize(A);
5       tq=0;  if ( t[1]<t[2],  tq=1; A=A~; );
6       t = matSVDcore(A);
```

```
 7       d = t[2];
 8       if ( tq,
 9           U = t[3];  V=t[1];
10       , /* else */
11           U = t[1];  V=t[3];
12       );
13       return( [U, d, V] );
14   }
```

For a numerical example see figure 28.5-A. The connection between the SVD of a matrix a and the eigenvectors of $A^T A$ is described in [158].

28.6 Goldschmidt's algorithm

A framework for the so-called *Goldschmidt algorithm* can be stated as follows. Initialize

$$x_0 \;=\; d^A, \qquad E_0 \;=\; d^B \tag{28.6-1a}$$

then iterate

$$P_{k+1} \;=\; 1 + \frac{1 - E_k}{a} \tag{28.6-1b}$$

$$x_{k+1} \;=\; x_k\, P_k^b \quad \to d^{A - B\,b/a} \tag{28.6-1c}$$

$$E_{k+1} \;=\; E_k\, P_k^a \quad \to 1 \tag{28.6-1d}$$

The algorithm converges quadratically. The updates for x and E (last two relations) can be computed independently. The iteration is not self-correcting, so the computations have to be carried out with full precision in all steps.

An invariant of the algorithm is given by x_k^a / E_k^b:

$$\frac{x_{k+1}^a}{E_{k+1}^b} \;=\; \frac{(x_k \cdot P_k^b)^a}{(E_k \cdot P_k^a)^b} \;=\; \frac{x_k^a}{E_k^b} \tag{28.6-2a}$$

Using

$$\frac{x_0^a}{E_0^b} \;=\; \frac{d^{A\,a}}{d^{B\,b}} \;=\; d^{A\,a - B\,b} \tag{28.6-2b}$$

and, as E converges to 1, we find that

$$\frac{x_\infty^a}{E_\infty^b} \;=\; x_\infty^a \;=\; \frac{x_0^a}{E_0^b} \tag{28.6-2c}$$

That is,

$$x_\infty \;=\; \left(\frac{x_0^a}{E_0^b}\right)^{1/a} \;=\; \frac{x_0}{E_0^{b/a}} \;=\; \frac{d^A}{d^{b\,B/a}} \;=\; d^{A - B\,b/a} \tag{28.6-2d}$$

One can now look for interesting special cases, b is set to one in what follows.

28.6.1 Algorithm for the a-th root

Solving $A - B/a = 1/a$ gives $B = A\,a - 1$ and especially $A = 0,\ B = 1$. That is, set

$$x_0 \;=\; d, \qquad E_0 \;=\; d^{a-1} \tag{28.6-3a}$$

then iterate

$$\begin{aligned} P_k &:= 1 + \frac{1 - E_k}{a} \quad \to 1 && \text{(28.6-3b)} \\ x_{k+1} &:= x_k \cdot P_k && \text{(28.6-3c)} \\ E_{k+1} &:= E_k \cdot P_k^a \quad \to 1 && \text{(28.6-3d)} \end{aligned}$$

until x close enough to $x_\infty = d^{\frac{1}{a}}$

Computation of the square root

For $a = 2$ one obtains an algorithm for the computation of the square root:

$$\sqrt{d} = d \prod_{k=0}^{\infty} \frac{3 - E_k}{2} \tag{28.6-4}$$

where $E_0 = d$, $E_{k+1} := E_k \, (\frac{3-E_k}{2})^2$.

28.6.2 Algorithm for the inverse a-th root

Solving $A - B/a = -1/a$ gives $B = A\,a + 1$ and especially $A = 1,\ B = a - 1$. That is, set

$$x_0 = 1, \qquad E_0 = d \tag{28.6-5}$$

then iterate as in formulas 28.6-3b..28.6-3d until x close enough to $x_\infty = d^{-\frac{1}{a}}$.

Computation of the inverse

Setting $a = 1$ one obtains an algorithm for the inverse ($P_k = 1 + (1 - E_k) = 2 - E_k$):

$$\frac{1}{d} = \prod_{k=0}^{\infty} (2 - E_k) \tag{28.6-6}$$

where $E_0 = d$, $E_{k+1} := E_k \, (2 - E_k)$.

Computation of the inverse square root

For $a = 2$ one obtains an algorithm for the inverse square root ($P_k = 1 + (1 - E_k)/2 = (3 - E_k)/2$):

$$\frac{1}{\sqrt{d}} = \prod_{k=0}^{\infty} \frac{3 - E_k}{2} \tag{28.6-7}$$

where $E_0 = d$, $E_{k+1} := E_k \, (\frac{3-E_k}{2})^2$.

28.6.3 Higher order algorithms for the inverse a-th root

Higher order iterations are obtained by appending higher terms to the expression $\left(1+\frac{1-E_k}{a}\right)$ in the definitions of P_{k+1} as suggested by equation 28.3-4 on page 565 and the identification $y = 1 - E$:

$$\begin{aligned} E_{k+1} &= E_k\, P_k^a \quad \text{where} && \text{(28.6-8)}\\ P_k &= 1 + \frac{1-E_k}{a} \qquad \text{[second order]} && \text{(28.6-9)}\\ &\quad + \frac{(1+a)\,(1-E_k)^2}{2\,a^2} \qquad \text{[third order]}\\ &\quad + \frac{(1+a)\,(1+2a)\,(1-E_k)^3}{6\,a^3} \qquad \text{[fourth order]}\\ &\quad + \ \ldots\ +\\ &\quad + \frac{(1+a)\,(1+2a)\ldots(1+(n-1)\,a)\,(1-E_k)^n}{n!\,a^n} \qquad \text{[order } (n+1)\text{]} \end{aligned}$$

$x_0 = 1$ $E_0 = 2.0$ $P_0 = 0.90625$ $b_0 = 0.0$
$x_1 = 0.90625$ $E_1 = 1.3490314483642578125$ $P_1 = 0.93177697415069360431516543030738830566640625$ $b_1 = 1.5185$
$x_2 = 0.84442288282406607891061867121607065200805664062 5$ $E_2 = 1.01688061936626457410433320872039209492188542219092968$ $P_2 = 0.99582436942565084187883150540345532399967189056293558 21$ $b_2 = 5.8884$
$x_3 = 0.84089688481686585202128466050063877105975287186278309 56$ $E_3 = 1.00000223363323283559583877024921007574068879685671957$ $P_3 = 0.99999944159247134069773211917093099758090030134706071 62$ $b_3 = 18.772$
$x_4 = 0.84089641525371454412926831199731186378490804857313364 97$ $E_4 = 1.00000000000000000522367709431971479704373188248422463 7$ $P_4 = 0.99999999999999999869408072642007130500262990214776549 07$ $b_4 = 57.409$
$x_5 = 0.84089641525371454303112547623321489504003426235678452 49$ $E_5 = 1.00000000000000000000000000000000000000000000000000000 67$ $P_5 = 0.99999999999999999999999999999999999999999999999999999 9833$ $b_5 = 173.32$
$1/\sqrt[4]{2} = 0.84089641525371454303112547623321489504003426235678451 08\ldots$

Figure 28.6-A: Numerical quantities occurring in the computation of $1/\sqrt[2]{2}$ using a third order Goldschmidt algorithm. The value of b_k gives the number of correct bits at step k.

As an example, the inverse fourth root of $d = 2$ can be obtained via the third order algorithm

$$\begin{aligned} x_0 &= 1 && \text{(28.6-10a)}\\ E_0 &= d\ =\ 2 && \text{(28.6-10b)}\\ E_{k+1} &= E_k\, P_k^4\ =\ E_k\,\left(\frac{45 - 18\,E_k + 5\,E_k^2}{32}\right)^4 && \text{(28.6-10c)}\\ x_{k+1} &= x_k\, P_k && \text{(28.6-10d)} \end{aligned}$$

Figure 28.6-A shows the numerical values of x_k, E_k and P_k up to step $k = 6$. The approximate precision in bits of x_k is computed as $b_k = -\log(|1 - E_k|)/\log(2)$.

28.7 Products for the a-th root

Rewrite the well-known product form

$$\frac{1}{1-y} \quad = \quad (1+y)\,(1+y^2)\,(1+y^4)\,(1+y^8)\,\ldots \tag{28.7-1}$$

as

$$\frac{1}{1-y} \quad = \quad \prod_{k>0}(1+Y_k) \quad \text{where} \quad Y_1 \;:=\; y, \qquad Y_{k+1} \;:=\; Y_k^2 \tag{28.7-2}$$

We give product forms for a-th roots and their inverses that generalize the relations above.

28.7.1 Second order products

Products for the square root and its inverse

For the inverse square root use $1/\sqrt{1-y} = (1+y/2)\cdot 1/\sqrt{1-y^2/4\,(3+y)}$, thereby

$$\frac{1}{\sqrt{1-y}} \quad = \quad \prod_{k>0}(1+Y_k) \quad \text{where} \quad Y_1 \;:=\; \frac{y}{2}, \qquad Y_{k+1} \;:=\; Y_k^2\left(\frac{3}{2}+Y_k\right) \tag{28.7-3}$$

For the square root use $\sqrt{1-y} = (1-y/2)\cdot\sqrt{1-(y/(y-2))^2}$, so

$$\sqrt{1-y} \quad = \quad \prod_{k>0}(1+Y_k) \quad \text{where} \quad Y_1 \;:=\; -\frac{y}{2}, \qquad Y_{k+1} \;:=\; -\frac{1}{2}\left(\frac{Y_k}{1+Y_k}\right)^2 \tag{28.7-4}$$

Products for the a-th root and its inverse

The relation for the inverse a-th root is

$$\frac{1}{\sqrt[a]{1-y}} \quad = \quad (1-y)^{-1/a} \;=\; \prod_{k>0}(1+Y_k) \quad \text{where} \tag{28.7-5a}$$

$$Y_1 \quad := \quad \frac{y}{a}, \qquad Y_{k+1} \;:=\; \frac{1}{a}\left(1-(1-a\,Y_k)\,(Y_k+1)^a\right) \tag{28.7-5b}$$

Alternatively,

$$\frac{1}{\sqrt[a]{d}} \quad = \quad x\,(1-y)^{-1/a} \;=\; x\,\prod_{k>0}(1+Y_k) \tag{28.7-6}$$

with $y := 1 - d\,x^a$ and the definitions 28.7-5b for Y_k. For the a-th root we obtain

$$\sqrt[a]{1-y} \quad = \quad (1-y)^{1/a} \;=\; \prod_{k>0}(1+Y_k) \quad \text{where} \tag{28.7-7a}$$

$$Y_1 \quad := \quad \frac{y}{a}, \qquad Y_{k+1} \;:=\; \frac{1}{a}\,\frac{(1+a\,Y_k)-(1+Y_k)^a}{(1+Y_k)^a} \tag{28.7-7b}$$

28.7.2 Products of arbitrary order

We want to obtain an n-th order product for the inverse a-th root

$$\frac{1}{\sqrt[a]{1-y}} \;=\; \prod_{k>0}\left(1+T(Y_k)\right) \quad \text{where} \quad Y_1 \;=\; y, \qquad Y_{k+1} \;=\; N(Y_k) \tag{28.7-8}$$

The functions T and N have to be determined. Set

$$[1-y]^{-\frac{1}{a}} \;=\; \left(1+T(Y_1)\right)\,\left[\left(1+T(Y_1)\right)^a\,(1-y)\right]^{-\frac{1}{a}} \tag{28.7-9a}$$

$$=:\; \left(1+T(Y_1)\right)\,[1-Y_2]^{-\frac{1}{a}} \tag{28.7-9b}$$

where $1+T(Y_1)$ is the Taylor expansion

$$[1-y]^{-\frac{1}{a}} \;=\; 1+\frac{y}{a}+\frac{(1+a)\,y^2}{2\,a^2}+\frac{(1+a)(1+2a)\,y^3}{6\,a^3}+\ldots+\frac{\prod_{k=1}^{n-1}(1+k\,a)}{n!\,a^n}\,y^n+\ldots \tag{28.7-10}$$

up to order $n-1$. The Taylor expansion of Y_2 starts with a term $\sim y^n$. Using $Y_{k+1}=N(Y_k)$ as suggested by the relation between Y_2 and Y_1 gives a product with n-th order convergence. For example, for a third order product for $1/\sqrt{1-y}$, set

$$T(Y) \;:=\; \frac{1}{2}\,Y+\frac{3}{8}\,Y^2 \tag{28.7-11}$$

Now solve $\left(1+T(Y_1)\right)^2\,(1-y)=(1-Y_2)$ for Y_2 to obtain

$$Y_2 \;=\; \frac{5\,y^3}{8}+\frac{15\,y^4}{64}+\frac{9\,y^5}{64} \;=:\; N(y) \tag{28.7-12}$$

Then, finally, ($Y_1:=y$ and)

$$\frac{1}{\sqrt{1-y}} \;=\; \prod_{k>0}\left(1+T(Y_k)\right) \quad \text{where} \tag{28.7-13a}$$

$$T(Y) \;:=\; \frac{1}{2}\,Y+\frac{3}{8}\,Y^2 \tag{28.7-13b}$$

$$Y_{k+1} \;=\; N(Y_k) \;:=\; \frac{Y_k^3}{64}\,\left(40+15\,Y_k+9\,Y_k^2\right) \tag{28.7-13c}$$

Replace relation 28.7-13c by $Y_{k+1}=1-\left(1+T(Y_k)\right)^a\,(1-Y_k)$ to obtain the general formula for the inverse a-th root.

The second order products lead to expressions that are especially nice:

$$\frac{1}{\sqrt[a]{1-y}} \;=\; \prod_{k>0}\left(1+T(Y_k)\right) \quad \text{where} \quad T(y) \;:=\; +\frac{y}{a} \qquad \text{and} \tag{28.7-14a}$$

$$Y_{k+1} \;=\; N(Y_k) \;:=\; 1-\left(1+\frac{y}{a}\right)^a\,(1-y) \tag{28.7-14b}$$

$$\frac{1}{\sqrt[a]{1+y}} \;=\; \prod_{k>0}\left(1+T(Y_k)\right) \quad \text{where} \qquad T(y) \;:=\; -\frac{y}{a} \qquad \text{and} \tag{28.7-15a}$$

$$Y_{k+1} \;=\; N(Y_k) \;:=\; \left(1-\frac{y}{a}\right)^a\,(1+y)-1 \tag{28.7-15b}$$

$$\sqrt[a]{1-y} \;=\; \prod_{k>0}\left(1+T(Y_k)\right) \quad \text{where} \quad T(y) \;:=\; -\frac{y}{a} \qquad \text{and} \tag{28.7-16a}$$

$$Y_{k+1} \;=\; N(Y_k) \;:=\; \frac{(a-Y_k)^a-(1-Y_k)\,a^a}{(a-Y_k)^a} \;=\; 1-\frac{(1-Y_k)\,a^a}{(a-Y_k)^a} \tag{28.7-16b}$$

$$\sqrt[a]{1+y} \;=\; \prod_{k>0}\,(1+T(Y_k)) \quad \text{where} \quad T(y) \;:=\; +\frac{y}{a} \qquad \text{and} \tag{28.7-17a}$$

$$Y_{k+1} \;=\; N(Y_k) \;:=\; \frac{(1+Y_k)\,a^a-(a+Y_k)^a}{(a+Y_k)^a} \;=\; \frac{(1+Y_k)\,a^a}{(a+Y_k)^a}-1 \tag{28.7-17b}$$

The third order product for $\frac{1}{\sqrt[a]{1-y}}$ is

$$\frac{1}{\sqrt[a]{1-y}} \;=\; \prod_{k>0}\,(1+T(Y_k)) \quad \text{where} \quad T(y) \;=\; +\frac{y}{a}+\frac{y^2\,(1+a)}{2\,a^2} \qquad \text{and} \tag{28.7-18a}$$

$$Y_{k+1} \;=\; N(Y_k) \;:=\; 1-\left(1+\frac{y}{a}+\frac{y^2\,(1+a)}{2\,a^2}\right)^a\,(1-y) \tag{28.7-18b}$$

$$\;=\; 1-(1+T(y))^a\,(1-y) \tag{28.7-18c}$$

28.7.3 Third order product for the a-th root

```
? default(realprecision,55);
? a=3;d=2;al=a-1;be=a+1;
? F(x)=(al*x^a+be*d)/(be*x^a+al*d)  \\ == (x^3 + 4)/(2*x^3 + 2)
? p=99.0;  \\ very bad approximation to the root
? for(k=0,25,p*=F(p);print("  ",p););
  49.500153045449860867772853756570132948572606410388539 63
  24.750688696322535793295398076760659038190058852964934 60
  12.377792780128382599987309228382889563737720225044012 95
  6.1986817294679733089803948935359831901917837320165830 57
  3.1382160956525775164587135127175212698603091937687217 90
  1.7166434303974992394555143292360395395658208028534524 86
  1.2833235143227843328303771164010152645995928582804900 32
  1.2599262845711532794913599247538658268681639208667671 05
  1.2599210498948732250077509793665642207537327503965189 86
  1.2599210498948731647672106072782283505702514647015997 90
  1.2599210498948731647672106072782283505702514647015079 80
  1.2599210498948731647672106072782283505702514647015079 80
```

Figure 28.7-A: Computation of $\sqrt[3]{2}$ with a very bad initial approximation.

The third order iteration given as relation 28.2-9 on page 563 gives a simple product for $\sqrt[a]{d}$. Let

$$P_k \;:=\; \prod_{j=0}^{k} Y_k \tag{28.7-19a}$$

where Y_0 is sufficiently near to $\sqrt[a]{d}$ and $Y_k = F(P_{k-1})$ where

$$F(x) \;:=\; \frac{\alpha\,x^a+\beta\,d}{\beta\,x^a+\alpha\,d} \tag{28.7-19b}$$

with $\alpha = a-1$ and $\beta = a+1$. Then $P_\infty = \sqrt[a]{d}$. Figure 28.7-A shows the numerical quantities with the computation of $\sqrt[3]{2}$ with a starting value $Y_0 = 99$ that is not at all close to the root. We have $F(x) = \frac{x^3+4}{2x^3+2}$ which is $\approx \frac{1}{2}$ for large values of x. Therefore the big initial values are repeatedly halved before the third order convergence begins.

28.8 Divisionless iterations for polynomial roots

Let $f(x)$ be a polynomial in x with simple roots only, then

$$\Phi(x) \;:=\; x-p(x)\,f(x) \quad \text{where} \quad p(x) := f'(x)^{-1} \bmod f(x) \tag{28.8-1}$$

is a second order iteration for the roots of $f(x)$. The iteration involves no long division if all coefficients are small rationals. Instead of dividing by $f'(x)$ a multiplication by the modular inverse $p(x)$ is used. As $\deg(p) < \deg(f)$ we have $\deg(\Phi) \leq 2\,\deg(f) - 1$.

For example, for $f(x) = a\,x^2 + b\,x + c$ we obtain

$$\Phi(x) \quad = \quad x - \frac{2\,a\,x + b}{\Delta}\,f(x) \quad \text{where} \quad \Delta = b^2 - 4\,a\,c \tag{28.8-2}$$

The general expressions for polynomials of orders > 2 get complicated. However, for fixed polynomial coefficients the iteration is more manageable. For example, with $f(x) = x^3 + 5\,x + 1$ we obtain

$$\Phi(x) \quad = \quad x + \frac{f(x)}{527}\,\left(30\,x^2 - 9\,x + 100\right) \tag{28.8-3}$$

For the polynomial $x^n - d$ we have $p = x/(n\,d)$ and the iteration is (relation 28.3-16 on page 566):

$$\Phi(x) \quad = \quad x - \frac{x}{n\,d}\,(x^n - d) \;=\; x + \frac{1}{n}\,\left(x - \frac{x^{n+1}}{d}\right) \tag{28.8-4}$$

The construction is given in [148] where a method to construct divisionless iterations of arbitrary order is given: let $p\,f' + q\,f \equiv 1$, and

$$\Phi_1 \quad := \quad x - p_1\,f, \qquad p_1 \;:=\; p \tag{28.8-5a}$$

$$p_r \quad := \quad p\,p'_{r-1} - (r-1)\,q\,p_{r-1} \tag{28.8-5b}$$

$$\Phi_r \quad := \quad \Phi_{r-1} + (-1)^r\,p_r\,f^r/r! \tag{28.8-5c}$$

then Φ_r is an iteration of order $r + 1$.

Chapter 29

Iterations for the inversion of a function

In this chapter we study some general expressions for iterations for the zero of a function. Two schemes that give (arbitrary order, one-point) iterations, Householder's formula and Schröder's formula, are given. Several methods to construct alternative iterations are described. Moreover, iterations that also converge for multiple roots and a technique to turn a linear iteration into a super-linear one are presented.

29.1 Iterations and their rate of convergence

An *iteration* for a zero r (or root, $f(r) = 0$) of a function $f(x)$ are themselves functions $\Phi(x)$ that, when used like

$$x_{k+1} \quad = \quad \Phi(x_k) \tag{29.1-1}$$

will make x_k converge towards the root: $x_\infty = r$. Convergence is subject to the condition that x_0 was chosen close enough to r. The function $\Phi(x)$ must (and can) be constructed so that it has an attracting fixed point where $f(x)$ has a zero:

$$\Phi(r) \quad = \quad r \qquad \text{(fixed point)} \tag{29.1-2}$$

$$|\Phi'(r)| \quad < \quad 1 \qquad \text{(attracting)} \tag{29.1-3}$$

This type of iteration is called a *one-point iteration*. There are also *multi-point iterations*, these are of the form $x_{k+1} = \Phi(x_k, x_{k-1}, \ldots, x_{k-j}), j \geq 1$. An example is the two-point iteration known as the *secant method*

$$x_{k+1} \quad = \quad \Phi(x_k, x_{k-1}) \;=\; x_k - f(x_k)\,\frac{x_k - x_{k-1}}{f(x_k) - f(x_{k-1})} \tag{29.1-4}$$

We are mainly concerned with one-point iterations in what follows.

The order of convergence (or simply *order*) of a given iteration can be defined as follows: let $x = r \cdot (1+e)$ with $|e| \ll 1$ and $\Phi(x) = r \cdot (1 + \alpha e^n + O(e^{n+1}))$, then the iteration Φ is called *linear* (or first order) if $n = 1$ (and $|\alpha| < 1$). A linear iteration improves the result by (roughly) adding a constant amount of correct digits with every step.

A *super-linear* iteration does better than that: The number of correct digits grows exponentially (to the base n) at each step. Super-linear convergence of order n should better be called *exponential of order n*.

Iterations of second order ($n = 2$) are often called *quadratic* (or *quadratically convergent*), those of third order *cubic* iterations. Fourth, fifth and sixth order iterations are called *quartic*, *quintic* and *sextic* and so on. We note that the two-point iteration relation 29.1-4 has order $(\sqrt{5}+1)/2 \approx 1.618$, see [149, p.152].

It is conceivable to find iterations that do converge better than linear but less than exponential to any base: imagine an iteration that produces proportional k^2 digits at step k (this is *not* quadratical convergence which produces proportional 2^k correct digits at step k). That case is not covered by the 'order-n' notion just introduced. However, those (super-linear but sub-exponential) iterations are not usually encountered and we actually won't meet one. In fact, the constructions used in this chapter cannot produce such an iteration. For a more fine-grained definition of the concept of order see [63, p.21].

For $n \geq 2$ the iteration function Φ has a *super-attracting fixed point* at r: $\Phi'(r) = 0$. For an iteration of order n one has

$$\Phi'(r) = 0, \quad \Phi''(r) = 0, \quad \ldots, \quad \Phi^{(n-1)}(r) = 0 \tag{29.1-5}$$

There is no standard term for emphasizing the number of derivatives vanishing at the fixed point: *super-attracting of order n* might be appropriate.

To any iteration of order n for a function f one can add a term $f(x)^n \cdot \varphi(x)$ (where $\varphi(x)$ is an arbitrary function that is analytic in a neighborhood of the root) without changing the order of convergence. That term is assumed to be zero in what follows. The statement can easily be checked by verifying that the first $n-1$ derivatives of $\Phi_n(x) + f(x)^n \cdot \varphi(x)$, evaluated at the root r, equal zero.

Any two one-point iterations of the same order n differ by a term $f(x)^n \cdot \varphi(x)$.

Any two iterations of the same order n differ by a term $(x-r)^n \, \nu(x)$ where $\nu(x)$ is a function that is finite at r [149, p.174, ex.3].

Any one-point iteration of order n must explicitly evaluate f, f', ..., $f^{(n-1)}$ [255, p.98]. For methods to find zeros and extrema without evaluating derivatives see [63].

29.2 Schröder's formula

For $n \geq 2$ the expression

$$S_n(x) \;:=\; x + \sum_{k=1}^{n-1} (-1)^k \, \frac{f(x)^k}{k!} \left(\frac{1}{f'(x)} \frac{\partial}{\partial x} \right)^{k-1} \frac{1}{f'(x)} \tag{29.2-1}$$

gives a n-th order iteration for a (simple) root r of f [234, p.13]. That is,

$$\begin{aligned} S \;:=\; S_\infty(x) \;=\; x \;-\; & \frac{f}{1!\, f'} \;-\; \frac{f^2}{2!\, f'^3} \cdot f'' \;-\; \frac{f^3}{3!\, f'^5} \cdot \left(3f''^2 - f' f'''\right) \\ & - \frac{f^4}{4!\, f'^7} \cdot \left(15 f''^3 - 10 f' f'' f''' + f'^2 f''''\right) \\ & - \frac{f^5}{5!\, f'^9} \cdot \left(105 f''^4 - 105 f' f''^2 f''' + 10 f'^2 f'''^2 + 15 f'^2 f'' f'''' - f'^3 f'''''\right) - \ldots \end{aligned} \tag{29.2-2}$$

The second order iteration is the well-known Newton iteration. The third order iteration, obtained upon truncation after the third term on the right hand side, and written as

$$S_3 \;=\; x - \frac{f}{f'} \left(1 + \frac{f f''}{2 f'^2} \right) \tag{29.2-3}$$

is sometimes referred to as *Householder's method*. Approximating the second term on the right hand side gives *Halley's formula*:

$$H_3 \quad = \quad x - \frac{f}{f'}\left(1 - \frac{ff''}{2f'^2}\right)^{-1} \tag{29.2-4}$$

Write

$$S \quad = \quad x - U_1\frac{f}{1!\,f'} - U_2\frac{f^2}{2!\,f'^3} - U_3\frac{f^3}{3!\,f'^5} - \ldots - U_n\frac{f^n}{n!\,f'^{2n-1}} - \ldots \tag{29.2-5}$$

then $U_1 = 1$, $U_2 = f''$, $U_3 = 3f''^2 - f'f'''$, and we have the recursion

$$U_n \quad = \quad (2n-3)f''U_{n-1} - f'U'_{n-1} \tag{29.2-6}$$

An alternative recursion is given in [255, p.83], write

$$S \quad = \quad x - Y_1\left(\frac{f}{f'}\right) - Y_2\left(\frac{f}{f'}\right)^2 - Y_3\left(\frac{f}{f'}\right)^3 - \ldots - Y_n\left(\frac{f}{f'}\right)^n - \ldots \tag{29.2-7}$$

then $Y_1 = 1$ and

$$Y_n \quad = \quad \frac{1}{n}\left(2\,(n-1)\,\frac{f''}{2f'}\,Y_{n-1} - Y'_{n-1}\right) \tag{29.2-8}$$

Relation 29.2-1 on the preceding page with $f(x) = 1/x^a - d$ gives the division-free iteration 28.3-5 on page 565 for arbitrary order. For $f(x) = \log(x) - d$ one obtains the iteration 31.2-9 on page 623. For $f(x) = x^2 - d$ one obtains

$$S(x) \quad = \quad x - \left(\frac{x^2-d}{2x} + \frac{\left(x^2-d\right)^2}{8x^3} + \frac{\left(x^2-d\right)^3}{16x^5} + \frac{5\left(x^2-d\right)^4}{128x^7} + \ldots\right) \tag{29.2-9a}$$

$$= \quad x - 2x\cdot\left(Y + Y^2 + 2\,Y^3 + 5\,Y^4 + 14\,Y^5 + 42\,Y^6 + \ldots\right) \quad \text{where} \quad Y := \frac{x^2-d}{(2x)^2} \tag{29.2-9b}$$

The coefficients of the powers of Y are the Catalan numbers, see section 13.3 on page 320.

29.3 Schröder's formula and series reversion

We give three ways of obtaining Schröers iteration via (implicit or explicit) power series reversion. The *reversion* of a series

$$A(x) \quad = \quad \sum_{k=1}^{\infty} a_k\,x^k \tag{29.3-1}$$

is the series

$$B(x) \quad = \quad \sum_{k=1}^{\infty} b_k\,x^k \tag{29.3-2}$$

such that $A(B(x)) = x$. That is, $B(x) = A^{[-1]}(x)$ is the functional inverse of $A(x)$. Note that $A^{[-1]}(x)$ is *not* the same as $A^{-1}(x) = 1/A(x)$. A useful relation is given in [15, p.634]:

$$A^{[-1]}(x) \quad = \quad \sum_{k=1}^{\infty}\frac{x^k}{k!}\left[\left(\frac{\partial}{\partial x}\right)^{k-1}\left(\frac{x}{f(x)}\right)^k\right]_{x=0} \tag{29.3-3}$$

Equivalently, from [168, p.527],

$$k\,b_k \;=\; [k-1]\left(\frac{x}{A(x)}\right)^k \tag{29.3-4}$$

where $[k-1]\,Q$ is the $(k-1)$-st series coefficient of Q. We use the expression to give a few terms of the reversed series explicitly:

```
? n=5; R=O(x^(n+1));
? A=sum(k=1,n,x^k*eval(Str("a"k)))+R
  a1*x + a2*x^2 + a3*x^3 + a4*x^4 + a5*x^5 + O(x^6)
? B=sum(k=1,n,x^k/k*(polcoeff(truncate((x/A)^k),k-1)))+R
 + 1/a1 * x
 + (-a2/a1^3) * x^2
 + ((-a3*a1 + 2*a2^2)/a1^5) * x^3
 + ((-a4*a1^2 + 5*a3*a2*a1 - 5*a2^3)/a1^7) * x^4
 + ((-a5*a1^3 + (6*a4*a2 + 3*a3^2)*a1^2 - 21*a3*a2^2*a1 + 14*a2^4)/a1^9) * x^5 + O(x^6)
```

The same result is obtained with pari/gp's builtin function `serreverse()`. Relation 29.3-4 can be generalized for the n-th power of the reversed series:

$$[k]\;(B(x))^n \;=\; \frac{n}{k}\,[k-n]\left(\frac{x}{A(x)}\right)^k \tag{29.3-5}$$

This is the *Lagrange inversion formula*.

29.3.1 Method of deriving the inverse function

The starting point is the Taylor series of a function f around x_0:

$$f(x) \;=\; \sum_{k=0}^{\infty}\frac{1}{k!}f^{(k)}(x_0)\,(x-x_0)^k \tag{29.3-6a}$$

$$\;=\; f(x_0)+f'(x_0)\,(x-x_0)+\frac{1}{2}f''(x_0)\,(x-x_0)^2+\frac{1}{6}f'''(x_0)\,(x-x_0)^3+\ldots \tag{29.3-6b}$$

Now let $f(x_0)=y_0$ and r be the zero of f (that is, $f(r)=0$). We expand the inverse function $g=f^{-1}$ around y_0:

$$g(0) \;=\; \sum_{k=0}^{\infty}\frac{1}{k!}g^{(k)}(y_0)\,(0-y_0)^k \tag{29.3-7a}$$

$$\;=\; g(y_0)+g'(y_0)\,(0-y_0)+\frac{1}{2}g''(y_0)\,(0-y_0)^2+\frac{1}{6}g'''(y_0)\,(0-y_0)^3+\ldots \tag{29.3-7b}$$

Using $x_0=g(y_0)$ and $g(0)=r$ we obtain

$$r \;=\; x_0-g'(y_0)\,f(x_0)+\frac{1}{2}g''(y_0)\,f(x_0)^2-\frac{1}{6}g'''(y_0)\,f(x_0)^3+\ldots \tag{29.3-8}$$

Remains to express the derivatives of the inverse g in terms of (derivatives of) f. Set

$$f\circ g \;=\; \mathrm{id}\,,\qquad \text{that is:}\quad f(g(x)) \;=\; x \tag{29.3-9}$$

and derive the equation (chain rule) to obtain $g'(f(x))\,f'(x)=1$, so $g'(y)=\frac{1}{f'(x)}$. Derive $f(g(x))-x$ multiple times to obtain (arguments y of g and x of f are omitted for readability):

$$1 \;=\; f'g' \tag{29.3-10a}$$

$$0 \;=\; g'f''+f'^2g'' \tag{29.3-10b}$$

$$0 \;=\; g'f'''+3f'f''g''+f'^3g''' \tag{29.3-10c}$$

$$0 \;=\; g'f''''+4f'g''f'''+3f''^2g''+6f'^2f''g'''+f'^4g'''' \tag{29.3-10d}$$

This system of linear equations in the derivatives of g can be solved successively for g', g'', g''', ...:

$$g' = \frac{1}{f'} \tag{29.3-11a}$$

$$g'' = -\frac{f''}{f'^3} \tag{29.3-11b}$$

$$g''' = \frac{1}{f'^5}\left(3f''^2 - f'f'''\right) \tag{29.3-11c}$$

$$g'''' = \frac{1}{f'^7}\left(10f'f''f''' - 15f''^3 - f'^2f''''\right) \tag{29.3-11d}$$

$$g''''' = \frac{1}{f'^9}\left(105f''^4 - f'^3f''''' - 105f'f''^2f''' + 15f'^2f''f'''' + 10f'^2f'''^2\right) \tag{29.3-11e}$$

Thereby equation 29.3-8 can be written as

$$\begin{aligned} r &= x - \frac{1}{f'}f + \frac{1}{2}\left(-\frac{f''}{f'^3}\right)f^2 - \frac{1}{6}\left(\frac{1}{f'^5}\left(3f''^2 - f'f'''\right)\right)f^3 + \ldots \\ &= x - \frac{f}{1!\,f'} - \frac{f^2}{2!\,f'^3}\cdot f'' - \frac{f^3}{3!\,f'^5}\cdot\left(3f''^2 - f'f'''\right) - \ldots \end{aligned} \tag{29.3-12}$$

which is Schröder's iteration, equation 29.2-2 on page 582.

29.3.2 Method of reversing power series

Schröder's formula can be obtained as the reversion of the series

$$E(W) := -\sum_{k=1}^{\infty}\frac{f^{(k)}}{k!}(-W)^k = [1 - \exp(-W\,\partial)]\circ f \tag{29.3-13}$$

Let $L(W)$ be the reversion of $E(W)$, then $x - L(f)$ is Schröder's iteration:

```
1   ? n=4;  \\ up to order n
2   ? x=W;  \\ kludge for pari/gp's variable ordering
3   ? E=-sum(k=1,n,eval(Str("f"k))*(-x)^k/k!)+O(x^(n+1))
4     f1*W - 1/2*f2*W^2 + 1/6*f3*W^3 - 1/24*f4*W^4 + O(W^5)
5   ? L=serreverse(E)
6     1/f1*W + f2/(2*f1^3)*W^2 + ((-f3*f1 + 3*f2^2)/(6*f1^5))*W^3 + \
7     ((f4*f1^2 - 10*f3*f2*f1 + 15*f2^3)/(24*f1^7))*W^4 + O(W^5)
8   ? L=truncate(L);  \\ make it a polynomial
9   ? for(j=1,n,print(-polcoeff(L,j,x)*(f0)^j))
10    -f0/f1
11    -f0^2*f2/(2*f1^3)
12    (f0^3*f3*f1 - 3*f0^3*f2^2)/(6*f1^5)
13    (-f0^4*f4*f1^2 + 10*f0^4*f3*f2*f1 - 15*f0^4*f2^3)/(24*f1^7)
```

Here we use '`fk`' as a symbol for the k-th derivative of f.

29.3.3 Method of writing power series as operator functions

Write the power series of the function f symbolically as

$$T(f) = \sum_{k=0}^{\infty}\frac{f^{(k)}\,x^k}{k!} = \left[\exp\left(+h\,\frac{\partial}{\partial x}\right)\circ f\right]_{h=x} \tag{29.3-14}$$

In this notation Schröders formula becomes

$$S(x) = \left[\exp\left(-h\,\frac{\partial}{\partial f}\right)\circ x\right]_{h=f} \tag{29.3-15}$$

First expand as a series

$$S(x) \;=\; \left[\sum_{k=0}^{\infty} \frac{\left(-h\,\frac{\partial}{\partial f}\right)^k}{k!} \circ x\right]_{h=f} \;=\; \sum_{k=0}^{\infty} \frac{(-f)^k}{k!}\,\left[\left(\frac{\partial}{\partial f}\right)^k\, x\right] \tag{29.3-16}$$

Now use

$$\frac{\partial}{\partial f} \;=\; \frac{\partial x}{\partial f}\,\frac{\partial}{\partial x} \;=\; \frac{1}{f'}\,\frac{\partial}{\partial x} \tag{29.3-17}$$

and separate the term for $k = 0$ to obtain

$$S(x) \;=\; x + \sum_{k=1}^{\infty} \frac{(-f)^k}{k!}\,\left[\left(\frac{1}{f'}\frac{\partial}{\partial x}\right)^k\, x\right] \;=\; x + \sum_{k=1}^{\infty} \frac{(-f)^k}{k!}\,\left[\left(\frac{1}{f'}\frac{\partial}{\partial x}\right)^{k-1}\, \frac{1}{f'}\right] \tag{29.3-18}$$

Truncation gives relation 29.2-1.

29.4 Householder's formula

The following expression gives for $n \geq 2$ a n-th order iteration for a (simple) root r of f [149, p.169]:

$$H_n(x) \;:=\; x + (n-1)\,\frac{\left(\frac{1}{f(x)}\right)^{(n-2)}}{\left(\frac{1}{f(x)}\right)^{(n-1)}} \tag{29.4-1}$$

We refer to iterations of this type as *Householder* iterations, the name *König iteration function* is used in [261]. We have

$$H_2 \;=\; x - \frac{f}{f'} \tag{29.4-2a}$$

$$H_3 \;=\; x - \frac{2ff'}{2f'^2 - ff''} \tag{29.4-2b}$$

$$H_4 \;=\; x - \frac{3f(ff'' - 2f'^2)}{6ff'f'' - 6f'^3 - f^2f'''} \tag{29.4-2c}$$

$$H_5 \;=\; x + \frac{4f\left(6f'^3 - 6ff'f'' + f^2f'''\right)}{f^3f'''' - 24f'^4 + 36ff'^2f'' - 8f^2f'f''' - 6f^2f''^2} \tag{29.4-2d}$$

The second order variant is Newton's formula, the third order iteration is Halley's formula.

Kalantari and Gerlach [156] define the iteration

$$B_m = x - f\,\frac{D_{m-2}}{D_{m-1}} \tag{29.4-3a}$$

where $m \geq 2$, $D_0 = 1$, $D_1 = f'$, and

$$D_m \;=\; \det\begin{pmatrix} f' & \frac{f''}{2!} & \cdots & \frac{f^{(m-1)}}{(m-1)!} & \frac{f^{(m)}}{m!} \\ f & f' & \ddots & \ddots & \frac{f^{(m-1)}}{(m-1)!} \\ 0 & f & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \frac{f''}{2!} \\ 0 & 0 & \ddots & f & f' \end{pmatrix} \tag{29.4-3b}$$

The iteration turns out to be identical to the one of Householder ($B_n = H_n$). A recursive definition for D_m is given by

$$D_m = \sum_{i=1}^{m} (-1)^{i-1} f^{i-1} \frac{f^{(i)}}{i!} D_{m-i} \tag{29.4-4}$$

The derivation of Halley's formula by applying Newton's formula to $f/\sqrt{f'}$ can be generalized to produce m-order iterations as follows: Let $F_1 = f$ and for $m \geq 2$ let

$$F_m = \frac{F_{m-1}}{\sqrt[m]{F'_{m-1}}} \tag{29.4-5a}$$

$$H_m = x - \frac{F_{m-1}}{F'_{m-1}} \tag{29.4-5b}$$

That this recursion indeed gives the Householder iteration is shown in [156].

An alternative recursive formulation (also given in [156], ascribed to Ford/Pennline) is

$$Q_2 = 1 \tag{29.4-6a}$$

$$Q_m = f' Q_{m-1} - \frac{1}{m-2} f Q'_{m-1} \tag{29.4-6b}$$

$$H_m = x - f \frac{Q_m}{Q_{m-1}} \tag{29.4-6c}$$

The Taylor series of the k-th order Householder iteration around $f = 0$ up to order $k-1$ gives the k-th order Schröder iteration.

An *extraneous* fixed point of an iteration for a function f is a fixed point at z such that $f(z) \neq 0$. All extraneous fixed points for the iterations H_n are repelling ($|H(z)'| > 1$), see [157] and [261].

29.5 Dealing with multiple roots

The iterations given so far will not converge at the stated order if f has a multiple root at r. As an example consider the function

$$f(x) = (x^2 - d)^m \quad \text{where} \quad m \in \mathbb{N} \tag{29.5-1}$$

The iteration $\Phi(x) = x - f/f'$ is

$$\Phi_?(x) = x - \frac{x^2 - d}{m2x} \tag{29.5-2}$$

Its convergence is only linear for $m > 1$: $\Phi(\sqrt{d}(1+e)) = \sqrt{d}(1 + \frac{m-1}{m}e + O(e^2))$

Householder [149, p.161, ex.6] gives a second order iteration for a root of known multiplicity m as

$$\Phi_2(x) = x - m \cdot \frac{f}{f'} \tag{29.5-3}$$

Note that with the example above we obtain a quadratic iteration.

For roots of unknown multiplicity use the general expressions for iterations with $F := f/f'$ instead of f. Both F and f have the same set of roots, but all roots of F are simple. To illustrate this consider a function f that has a root of multiplicity m at r: $f(x) := (x-r)^m h(x)$ with $h(r) \neq 0$. Then

$$f'(x) = m\,(x-r)^{m-1} h(x) + (x-r)^m h'(x) \tag{29.5-4a}$$

$$= (x-r)^{m-1} \Big(m\,h(x) + (x-r)\,h'(x) \Big) \tag{29.5-4b}$$

and

$$F(x) = f(x)/f'(x) \quad = \quad (x-r)\,\frac{h(x)}{m\,h(x)+(x-r)\,h'(x)} \tag{29.5-5}$$

The fraction on the right hand side does not vanish at the root r.

Plugging $F = f/f'$ into Householder's formula (relation 29.4-1) we get the following iterations denoted by $H_k^{\%}$, the iterations H_k are given for comparison:

$$H_2 \quad = \quad x - \frac{f}{f'} \tag{29.5-6a}$$

$$H_2^{\%} \quad = \quad x - \frac{ff'}{f'^2 - ff''} \tag{29.5-6b}$$

$$H_3 \quad = \quad x - \frac{2ff'}{2f'^2 - ff''} \tag{29.5-6c}$$

$$H_3^{\%} \quad = \quad x + \frac{2f^2f'' - 2ff'^2}{2f'^3 - 3ff'f'' + f^2f'''} \tag{29.5-6d}$$

$$H_4 \quad = \quad x + \frac{3f^2f'' - 6ff'^2}{6f'^3 - 6ff'f'' + f^2f'''} \tag{29.5-6e}$$

$$H_4^{\%} \quad = \quad x + \frac{6ff'^3 + 3f^3f''' - 9f^2f'f''}{f^3f'''' - 6f'^4 + 12ff'^2f'' - 4f^2f'f''' - 3f^2f''^2} \tag{29.5-6f}$$

$$H_5 \quad = \quad x + \frac{24ff'^3 + 4f^3f''' - 24f^2f'f''}{f^3f'''' - 24f'^4 + 36ff'^2f'' - 8f^2f'f''' - 6f^2f''^2} \tag{29.5-6g}$$

The terms in the numerators and denominators of $H_k^{\%}$ and H_{k+1} are identical up to the integral constants. A alternative form for $H_k^{\%}$ is given by

$$H_k^{\%} \quad = \quad x + (k-1)\,\frac{(\log(f))^{(k-1)}}{(\log(f))^{(k)}} \tag{29.5-7}$$

Schröder's formula (relation 29.2-1), when inserting f/f', becomes:

$$\begin{aligned} S^{\%} \quad = \quad & x - \frac{ff'}{(f'^2 - ff'')} - \frac{f^2f'\left(ff'f''' - 2ff''^2 + f'^2f''\right)}{2\left(ff'' - f'^2\right)^3} - \\ & - \frac{f^3f'\left(2ff'^3f''f''' \pm \ldots - 3f^2f'^2f'''^2\right)}{6\left(ff'' - f'^2\right)^5} - \frac{f^4f'\left(3f'^8f'''' \pm \ldots - 36f^3f'^2f''^2f'''^2\right)}{24\left(ff'' - f'^2\right)^7} - \\ & - \ldots - \frac{f^kf'\left(\ldots\right)}{k!\left(ff'' - f'^2\right)^{2k-1}} \end{aligned} \tag{29.5-8}$$

We check the convergence with our example (relation 29.5-1), the second order iteration is

$$\Phi_2^{\%}(x) \quad = \quad S_2 \;=\; H_2 \;=\; x + x\frac{d-x^2}{d+x^2} = \frac{2\,d\,x}{x^2+d} \tag{29.5-9}$$

Convergence is indeed second order, as we have (compare to relation 28.2-7 on page 563)

$$\Phi_2^{\%}(\sqrt{d}\cdot\frac{1-e}{1+e}) \quad = \quad \sqrt{d}\cdot\frac{1-e^2}{1+e^2} \tag{29.5-10}$$

which holds independent of m. In general, with $H_k^{\%}$ one obtains for the square root:

$$\Phi_k^{\%}(\sqrt{d}\cdot\frac{1-e}{1+e}) \quad = \quad \sqrt{d}\cdot\frac{1-e^k}{1+e^k} \tag{29.5-11}$$

Using Schröder's third order formula for f/f' with f as in 29.5-1 we obtain a beautiful *fourth* order iteration for $\sqrt{d}$:

$$S_3^{\%}(x) = x + x\,\frac{d-x^2}{d+x^2} + x\,d\,\frac{(d-x^2)^2}{(d+x^2)^3} \tag{29.5-12a}$$

$$S_3^{\%}(\sqrt{d}\,\frac{1-e}{1+e}) = \sqrt{d}\,\frac{1+3e^2-3e^4-e^6}{1+3e^2+3e^4+e^6} \tag{29.5-12b}$$

$$= \sqrt{d}\,\frac{1-c}{1+c} \quad \text{where} \quad c = e^4\,\frac{e^2+3}{3e^2+1} \tag{29.5-12c}$$

In general, the $(1+a\,k)$-th order Schröder iteration for $1/\sqrt[a]{d}$ obtained through f/f' has an order of convergence that exceeds the expected order by one. The third order Schröder iteration for $f(x) = 1-d\,x^2$ is

$$S_3^{\%}(x) = x + x\,\frac{1-dx^2}{1+dx^2} + x\,\frac{(1-dx^2)^2}{(1+dx^2)^3} \tag{29.5-13}$$

The iteration also has 4th order convergence and the error expression $S_3^{\%}(\frac{1}{\sqrt{d}}\,\frac{1-e}{1+e})$ is the same (replacing $\sqrt{d}$ by $1/\sqrt{d}$) as in relation 29.5-12b.

29.6 More iterations

We give expressions for iterations via Padé approximants, radicals, and show how iterations can be obtained from given ones. Finally we give one form of a multi-point iteration.

29.6.1 Rational iterations from Padé approximants

The $[i,j]$-th Padé approximant of Φ_n in f gives an iteration of order $p = i+j+1$ (if $n \geq p$). Write $P_{[i,j]}$ for an iteration (of order $i+j+1$) that is obtained using the Padé approximant $[i,j]$. For the second order (where the Newton iteration is $P_{[1,0]}(x) = x - \frac{f}{f'}$) this method gives one alternative form, namely

$$P_{[0,1]}(x) = x^2\,\frac{f'}{f+x\,f'} = x - \frac{x\,f}{f+x\,f'} = x - \frac{x\,f}{(x\,f)'} = x\,\left(1+\frac{f}{x\,f'}\right)^{-1} \tag{29.6-1}$$

For the third order we find $P_{[2,0]}(x) = S_3(x)$, $P_{[1,1]}(x) = H_3(x)$, and

$$P_{[0,2]}(x) = \frac{2x^3 f'^3}{2f^2f' + 2xff'^2 + xf^2f'' + 2x^2f'^3} \tag{29.6-2a}$$

$$= x - \frac{xf\left(2ff' + xff'' + 2xf'^2\right)}{\left(2f^2f' + 2xff'^2 + xf^2f'' + 2x^2f'^3\right)} \tag{29.6-2b}$$

$$= x\,\left(1 + \frac{f}{xf'} + \frac{f^2}{x^2f'^2} + \frac{f^2f''}{2xf'^3}\right)^{-1} \tag{29.6-2c}$$

$$= x\,\left(1 + \frac{f}{(x\,f')} + \frac{f^2(x^2\,f'')}{2(x\,f')^3} + \frac{f^2}{(x\,f')^2}\right)^{-1} \tag{29.6-2d}$$

Alternatively one can use the Padé approximant $A_{[i,j]}$ of $(\Phi(x)-x)/f$ in f where $\Phi(x)$ is a given iteration

of order $\geq i+j+2$. Then $\Phi^+_{[i,j]} := x + f \cdot A_{[i,j]}$ is an iteration which has order $n = i+j+2$.

$$\Phi^+_{[0,1]}(x) \;=\; x - \frac{2ff'}{2f'^2 - ff''} \;=\; H_3(x) \tag{29.6-3a}$$

$$\Phi^+_{[1,0]}(x) \;=\; x - \frac{f\left(f\,f'' + 2\,f'^2\right)}{2\,f'^3} \;=\; S_3(x) \tag{29.6-3b}$$

The iterations Φ^+ of order n are expressions in $x, f, f', \ldots, f^{(n-1)}$. Fourth order iterations are

$$\Phi^+_{[2,0]}(x) \;=\; x - \frac{f\left(6\,f'^4 + 3\,f\,f'^2\,f'' - f^2\,f'\,f''' + 3\,f^2\,f''^2\right)}{6\,f'^5} \tag{29.6-4a}$$

$$\;=\; x - \frac{f}{f'} - \frac{f^2}{2\,f'^3}\cdot f'' - \frac{f^3}{6\,f'^5}\cdot\left(3f''^2 - f'\,f'''\right) \;=\; S_4(x) \tag{29.6-4b}$$

$$\;=\; x - \frac{f^2}{f'^2}\left(\frac{f'}{f} + \frac{f''}{2\,f'} + \frac{f\left(3\,f''^2 - f'\,f'''\right)}{6\,f'^3}\right) \tag{29.6-4c}$$

$$\Phi^+_{[1,1]}(x) \;=\; x - \frac{f\left(2\,f\,f'\,f''' - 3\,f\,f''^2 + 6\,f'^2\,f''\right)}{f'\left(2\,f\,f'\,f''' - 6\,f\,f''^2 + 6\,f'^2\,f''\right)} \tag{29.6-4d}$$

$$\Phi^+_{[0,2]}(x) \;=\; x - \frac{12\,f\,f'^3}{\left(12\,f'^4 - 6\,f\,f'^2\,f'' + 2\,f^2\,f'\,f''' - 3\,f^2\,f''^2\right)} \tag{29.6-4e}$$

$$\;=\; x - \left(\frac{f'}{f} - \frac{f''}{2\,f'} - \frac{f\left(3\,f''^2 - 2\,f'\,f'''\right)}{12\,f'^3}\right)^{-1} \tag{29.6-4f}$$

The iteration $\Phi^+_{[n,0]}$ always coincides with Schröder's iteration. In general one obtains $n-1$ additional forms of iterations using the approximants $[0, n-2]$, $[1, n-3]$, $\ldots$, $[n-3, 1]$.

Neglecting terms that contain the third derivative in relation 29.6-4d we obtain the third order iteration

$$\Phi_3 \;=\; x - \frac{f\left(2\,f'^2 - f\,f''\right)}{f'\left(2\,f'^2 - 2\,f\,f''\right)} \tag{29.6-5}$$

A closed form for the Padé approximants of the r-th root is given in relation 28.2-13 on page 564.

29.6.2 An iteration involving radicals

By directly solving the truncated Taylor expansion

$$f(r) \;=\; f(x) + f'(x)\,(r-x) + \frac{1}{2}f''(x)\,(r-x)^2 \tag{29.6-6}$$

of $f(r) = 0$ around x one obtains the following third order iteration:

$$\Phi_3 \;=\; x - \frac{1}{f''}\left(f' \pm \sqrt{f'^2 - 2ff''}\right) = x - \frac{f'}{f''}\left(1 \pm \sqrt{1 - 2\,\frac{f\,f''}{f'^2}}\right) \tag{29.6-7}$$

For $f(x) = a\,x^2 + b\,x + c$ it gives the two solutions of the quadratic equation $f(x) = 0$; for other functions one obtains an iterated square root expression for the roots.

The following form, given in [255, p.94], avoids possible cancellation:

$$\Phi_3 \quad = \quad x - \frac{2\,u}{1+\sqrt{1-4\,A\,u}} \quad \text{where} \quad u = f/f' \quad \text{and} \quad A = f''/(2f) \tag{29.6-8}$$

It can be obtained by observing that

$$\frac{-b+\sqrt{b^2-4ac}}{2a} \quad = \quad \frac{-2c}{b+\sqrt{b^2-4ac}} \tag{29.6-9}$$

29.6.3 Iterations from iterations

Alternative rational forms can also be obtained in a way that generalizes the method used for multiple roots: if we emphasize the so far notationally omitted dependency from the function f as $\Phi\{f\}$. The iteration $\Phi\{f\}$ has fixed points where f has a root r, so $x - \Phi\{f\}$ again has a root at r. Hence we can build more iterations that will converge to those roots as $\Phi\{x - \Phi\{f\}\}$. For dealing with multiple roots we used $\Phi\{x - \Phi_2\{f\}\}_k = \Phi\{f/f'\}$. An iteration $\Phi_k\{x - \Phi_j\{f\}\}$ can only be expected to have a k-th order convergence.

29.6.4 A multi-point iteration

A multi-point iteration of order r can be given [255, p.165] as

$$\Phi_r(x) \quad = \quad \Phi_{r-1}(x) - \frac{f\left(\Phi_{r-1}(x)\right)}{f'(x)} \tag{29.6-10}$$

where $\Phi_{r-1}(x)$ is an iteration of order $r-1$. For example, choose $\Phi_2(x) = x - f(x)/f'(x)$ to obtain the third order iteration

$$\Phi_3(x) \quad = \quad \Phi_2(x) - \frac{f\left(\Phi_2(x)\right)}{f'(x)} \tag{29.6-11}$$

Apply the method again to obtain

$$\Phi_4(x) \quad = \quad \Phi_3(x) - \frac{f\left(\Phi_3(x)\right)}{f'(x)} \quad = \quad \Phi_2(x) - \frac{f\left(\Phi_2(x)\right)}{f'(x)} - \frac{f\left(\Phi_3(x)\right)}{f'(x)} \tag{29.6-12}$$

The r-th order iteration is

$$\Phi_r(x) \quad = \quad \Phi_2(x) - \sum_{k=2}^{r-1} \frac{f\left(\Phi_k(x)\right)}{f'(x)} \tag{29.6-13a}$$

$$\quad = \quad x - \frac{1}{f'(x)} \left[f(x) + \sum_{k=2}^{r-1} f\left(\Phi_k(x)\right) \right] \tag{29.6-13b}$$

The function f is evaluated at $r-1$ points but the derivative is only evaluated at x. The iteration also involves only one inversion.

29.7 Convergence improvement by the delta squared process

Given a sequence of partial sums x_k the so-called *delta squared process* computes a new sequence x_k^* of extrapolated sums:

$$x_k^* \quad = \quad x_{k+2} - \frac{(x_{k+2}-x_{k+1})^2}{x_{k+2} - 2\,x_{k+1} + x_k} \tag{29.7-1}$$

The method is due to Aitken. The name 'delta squared process' is due to the fact that the formula has the form

$$x^* = x - \frac{(\Delta x)^2}{(\Delta^2 x)} \tag{29.7-2}$$

where Δ is the difference operator. Note that the mathematically equivalent form

$$x_k^* = \frac{x_k\, x_{k+2} - x_{k+1}^2}{x_{k+2} - 2\, x_{k+1} + x_k} \tag{29.7-3}$$

sometimes given should be avoided with numerical computations due to possible cancellation.

If $x_k = \sum_{i=0}^{k} a_i$ and the ratio of consecutive summands a_i is approximately constant (that is, a is close to a geometric series) then x^* converges significantly faster to x_∞ than x. Rewrite relation 29.7-1 with $a_k := x_k - x_{k-1}$:

$$x_k^* = x_{k+2} - \frac{(a_{k+2})^2}{a_{k+2} - a_{k+1}} \tag{29.7-4}$$

Then for a geometric series (where $a_{k+1}/a_k = q$) we have

$$x_k^* = x_{k+2} - \frac{(a_{k+2})^2}{a_{k+2} - a_{k+1}} = x_{k+2} - \frac{\left(a_0\, q^{k+2}\right)^2}{a_0\,(q^{k+2} - q^{k+1})} \tag{29.7-5a}$$

$$= a_0\, \frac{1 - q^{k+3}}{1 - q} + a_0\, q^{k+2} \cdot \frac{q^{k+2}}{q^{k+1} - q^{k+2}} = \frac{a_0}{1 - q}\left(1 - q^{k+3} + q^{k+3}\right) \tag{29.7-5b}$$

$$= \frac{a_0}{1 - q} \tag{29.7-5c}$$

which is the exact sum. Now consider the sequence

$$x_0, \quad x_1 = \Phi(x_0), \quad x_2 = \Phi(x_1) = \Phi(\Phi(x_0)), \quad \ldots \tag{29.7-6}$$

of successively better approximations to some root r of a function f. Think of the x_k as partial sums of a series whose sum is the root r. Apply the idea to define an improved iteration Φ^* from a given one Φ:

$$\Phi^*(x) = \Phi(\Phi(x)) - \frac{[\Phi(\Phi(x)) - \Phi(x)]^2}{\Phi(\Phi(x)) - 2\,\Phi(x) + x} \tag{29.7-7}$$

The good news is that $\Phi*$ will give quadratic convergence even if Φ only has linear convergence. As an example take $f(x) = (x^2 - d)^2$, forget that its root $\sqrt{d}$ is a double root, and happily define $\Phi(x) = x - f(x)/f'(x) = x - (x^2 - d)/(4x)$. Convergence is only linear:

$$\Phi(\sqrt{d}\cdot(1+e)) = \sqrt{d}\cdot\left(1 + \frac{e}{2} + \frac{e^2}{4} + O(e^3)\right) \tag{29.7-8}$$

Then try

$$\Phi^*(x) = \frac{d\,(7\,x^2 + d)}{x\,(3\,x^2 + 5\,d)} \tag{29.7-9}$$

and find that it has quadratic convergence

$$\Phi(\sqrt{d}\cdot(1+e)) = \sqrt{d}\cdot\left(1 - \frac{e^2}{4} + \frac{e^3}{16} + O(e^4)\right) \tag{29.7-10}$$

In general, if Φ_n has convergence of order $n > 1$ then Φ_n^* will be of order $2n - 1$, but linear convergence ($n = 1$) is turned into second order, see [149, p.165].

Chapter 30

The arithmetic-geometric mean (AGM)

The arithmetic-geometric mean (AGM) is the basis for fast algorithms for the computation of π to high precision. We give AGM based algorithms for the computation of certain hypergeometric functions. AGM-based algorithms for the computation of the logarithm are given in section 31.1.1 on page 617, and for the exponential function in section 31.2.1 on page 623.

30.1 The AGM

The *arithmetic-geometric mean* (AGM) plays a central role in the high precision computation of logarithms and π. The AGM(a, b) is defined as the limit of the iteration

$$a_{k+1} = \frac{a_k + b_k}{2} \tag{30.1-1a}$$

$$b_{k+1} = \sqrt{a_k\, b_k} \tag{30.1-1b}$$

starting with $a_0 = a$ and $b_0 = b$. Both of the values converge quadratically to a common limit. The related quantity c_k used in many AGM based computations is defined as

$$c_k^2 = a_k^2 - b_k^2 = (a_{k-1} - a_k)^2 \tag{30.1-2}$$

An alternative way for the computation for the AGM iteration is

$$c_{k+1} = \frac{a_k - b_k}{2} \tag{30.1-3a}$$

$$a_{k+1} = \frac{a_k + b_k}{2} \tag{30.1-3b}$$

$$b_{k+1} = \sqrt{a_{k+1}^2 - c_{k+1}^2} \tag{30.1-3c}$$

Schönhage gives the most economic variant of the AGM, which, apart from the square root, only needs

one squaring per step:

$$\begin{aligned}
A_0 &= a_0^2 && (30.1\text{-}4a)\\
B_0 &= b_0^2 && (30.1\text{-}4b)\\
t_0 &= 1-(A_0-B_0) && (30.1\text{-}4c)\\
S_k &= \frac{A_k+B_k}{4} && (30.1\text{-}4d)\\
b_k &= \sqrt{B_k} \qquad \text{[square root]} && (30.1\text{-}4e)\\
a_{k+1} &= \frac{a_k+b_k}{2} && (30.1\text{-}4f)\\
A_{k+1} &= a_{k+1}^2 \qquad \text{[squaring]} && (30.1\text{-}4g)\\
&= \left(\frac{\sqrt{A_k}+\sqrt{B_k}}{2}\right)^2 = \frac{A_k+B_k}{4}+\frac{\sqrt{A_k\,B_k}}{2} && (30.1\text{-}4h)\\
B_{k+1} &= 2\,(A_{k+1}-S_k) \;=\; b_{k+1}^2 && (30.1\text{-}4i)\\
c_{k+1}^2 &= A_{k+1}-B_{k+1} \;=\; a_{k+1}^2-b_{k+1}^2 && (30.1\text{-}4j)\\
t_{k+1} &= t_k-2^{k+1}\,c_{k+1}^2 && (30.1\text{-}4k)
\end{aligned}$$

Starting with $a_0=A_0=1$, $B_0=1/2$ one has $\pi\approx(2\,a_n^2)/t_n$. The importance of the AGM is related to the fact that it can be used to compute certain hypergeometric functions fast. Indeed, one has

$$F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\ 1\end{matrix}\,\middle|\,1-\frac{b^2}{a^2}\right) \;=\; \frac{a}{\mathrm{AGM}(a,b)} \;=\; \frac{1}{\mathrm{AGM}(1,b/a)} \tag{30.1-5}$$

The relation is usually written as

$$F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\ 1\end{matrix}\,\middle|\,k\right) \;=\; \frac{1}{\mathrm{AGM}(1,\sqrt{1-k})} \tag{30.1-6}$$

corresponding to $\mathrm{AGM}(1,k)$, that is

$$a_0=1, \qquad b_0=k, \quad \text{and} \quad c_0=\sqrt{1-k^2} \tag{30.1-7}$$

The quantity

$$R'(k) \;:=\; 1-\frac{1}{2}\sum_{n=0}^{\infty}2^n c_n^2 \tag{30.1-8}$$

together with the AGM leads to a fast algorithm for the function $F\left(\begin{matrix}-\frac{1}{2},\,\frac{1}{2}\\ 1\end{matrix}\,\middle|\,k\right)$, see section 30.2.

Combining two steps of the AGM iteration leads to the fourth order AGM iteration:

$$\begin{aligned}
\alpha_0 &= \sqrt{a_0} && (30.1\text{-}9a)\\
\beta_0 &= \sqrt{b_0} && (30.1\text{-}9b)\\
\alpha_{k+1} &= \frac{\alpha_k+\beta_k}{2} && (30.1\text{-}9c)\\
\beta_{k+1} &= \left(\frac{\alpha_k\,\beta_k\,(\alpha_k^2+\beta_k^2)}{2}\right)^{1/4} && (30.1\text{-}9d)\\
\gamma_k^4 &= \alpha_k^4-\beta_k^4 \qquad = c_{2k}^2 && (30.1\text{-}9e)
\end{aligned}$$

We have $\alpha_k = \sqrt{a_{2k}}$, $\beta_k = \sqrt{b_{2k}}$. An alternative formulation of the iteration is:

$$\gamma_{k+1} = \frac{\alpha_k - \beta_k}{2} \tag{30.1-10a}$$

$$\alpha_{k+1} = \frac{\alpha_k + \beta_k}{2} \tag{30.1-10b}$$

$$\beta_{k+1} = \left(\alpha_{k+1}^4 - \gamma_{k+1}^4\right)^{1/4} \tag{30.1-10c}$$

$$c_{2k}^2 + 2\,c_{2k+1}^2 = \alpha_{k-1}^4 - \left(\alpha_k^2 - \gamma_k^2\right)^2 \tag{30.1-10d}$$

Compute R' via

$$R'(k) = 1 - \sum_{n=0}^{\infty} 4^n \left(\alpha_n^4 - \left(\frac{\alpha_n^2 + \beta_n^2}{2}\right)^2\right) \tag{30.1-11}$$

30.2 The elliptic functions K and E

The elliptic functions $K(k)$ and $E(k)$ can be computed via the AGM which gives super-linear convergence. The logarithmic singularity of $K(k)$ at the point $k = 1$ (relation 30.2-4, see also relation 31.1-1a on page 617) is the key to the fast computation of the logarithm. The exponential function could be computed by inverting the logarithm but also as described in section 31.2.1 on page 623. For computations with very high precision the algorithms based on the elliptic functions are the fastest known today for the logarithm, the number π, and the exponential function.

30.2.1 Elliptic K

The function K can be defined as

$$K(k) = \int_0^{\pi/2} \frac{d\,\vartheta}{\sqrt{1 - k^2\,\sin^2\vartheta}} = \int_0^1 \frac{d\,t}{\sqrt{(1-t^2)\,(1-k^2t^2)}} \tag{30.2-1}$$

One has

$$K(k) = \frac{\pi}{2}\,F\left(\begin{matrix}\frac{1}{2},\frac{1}{2}\\ 1\end{matrix}\,\middle|\,k^2\right) \tag{30.2-2a}$$

$$= \frac{\pi}{2}\sum_{i=0}^{\infty}\left(\frac{(2\,i-1)!!}{2^i\,i!}\right)^2 k^{2i} = \frac{\pi}{2}\sum_{i=0}^{\infty}\left(\frac{\binom{2\,i}{i}}{4^i}\right)^2 k^{2i} \tag{30.2-2b}$$

$$= \frac{\pi}{2}\left(1 + \left(\frac{1}{2}\right)^2 k^2 + \left(\frac{1\cdot 3}{2\cdot 4}\right)^2 k^4 + \left(\frac{1\cdot 3\cdot 5}{2\cdot 4\cdot 6}\right)^2 k^6 + \ldots\right) \tag{30.2-2c}$$

$$= \frac{\pi}{2}\left(1 + \frac{1}{4}\,k^2 + \frac{9}{64}\,k^4 + \frac{25}{256}\,k^6 + \frac{1225}{16384}\,k^8 + \frac{3969}{65536}\,k^{10} + \ldots\right) \tag{30.2-2d}$$

See section 35.2.7 on page 700 for transformations in terms of hypergeometric functions. Special values are $K(0) = \frac{\pi}{2}$ and $\lim_{k\to 1-} K(k) = +\infty$.

The computational interesting form is $F\left(\begin{matrix}\frac{1}{2},\frac{1}{2}\\ 1\end{matrix}\,\middle|\,z\right) = 1/\operatorname{AGM}(1,\sqrt{1-z})$ (see section 30.1 on page 593):

$$K(k) = \frac{\pi}{2\,\operatorname{AGM}(1,k')} = \frac{\pi}{2\,\operatorname{AGM}(1,\sqrt{1-k^2})} \tag{30.2-3a}$$

One defines $k' = \sqrt{1-k^2}$ and $K'(k)$ as $K(k')$:

$$K'(k) := K(\sqrt{1-k^2}) = \frac{\pi}{2\,\operatorname{AGM}(1,k)} \tag{30.2-3b}$$

A C++ implementation of the AGM based computation is given in [hfloat: src/tz/elliptic-k.cc].

For k close to 1 we have

$$K(k) \quad \approx \quad \log \frac{4}{\sqrt{1-k^2}} \tag{30.2-4}$$

The following estimate is given in [56, p.11]:

$$\left| K'(k) - \log \frac{4}{k} \right| \quad \leq \quad 4\,k^2\,(8 + \log k) \quad \text{where} \quad 0 < k \leq 1 \tag{30.2-5}$$

Product forms for K and K' that are also candidates for fast computations are, for $0 < k_0 \leq 1$,

$$\frac{2}{\pi}\,K'(k_0) \quad = \quad \prod_{n=0}^{\infty} \frac{2}{1+k_n} \quad = \quad \prod_{n=1}^{\infty} 1 + k'_n \quad \text{where} \quad k_{n+1} \;:=\; \frac{2\sqrt{k_n}}{1+k_n}, \quad k_\infty = 1 \tag{30.2-6a}$$

$$\frac{2}{\pi}\,K'(k_0) \quad = \quad \prod_{n=0}^{\infty} \frac{1}{\sqrt{k_n}} \quad \text{where} \quad k_{n+1} \;:=\; \frac{1+k_n}{2\sqrt{k_n}}, \quad k_\infty = 1 \tag{30.2-6b}$$

The second form is computationally especially attractive since, apart from the multiplication with the main product, only a inverse square root needs to be computed per step. The product formulas follow directly from relation 30.2-3b (and $\mathrm{AGM}(a,b) = a\ \mathrm{AGM}(1,b/a) = b\ \mathrm{AGM}(a/b,1)$):

$$\frac{1}{\mathrm{AGM}(1,k)} \quad = \quad \left[\mathrm{AGM}\left(\frac{1+k}{2}, \sqrt{k}\right)\right]^{-1} \tag{30.2-7a}$$

$$= \quad \left[\frac{1+k}{2}\ \mathrm{AGM}\left(1, \frac{2\sqrt{k}}{1+k}\right)\right]^{-1} \qquad \text{(first form)} \tag{30.2-7b}$$

$$= \quad \left[\sqrt{k}\ \mathrm{AGM}\left(\frac{1+k}{2\sqrt{k}}, 1\right)\right]^{-1} \qquad \text{(second form)} \tag{30.2-7c}$$

Similarly, for $0 < k_0 \leq 1$,

$$\frac{2}{\pi}\,K(k_0) \quad = \quad \prod_{n=0}^{\infty} \frac{2}{1+k'_n} \quad = \quad \prod_{n=1}^{\infty} 1 + k_n \quad \text{where} \quad k_{n+1} \;:=\; \frac{1-k'_n}{1+k'_n}, \quad k_\infty = 0 \tag{30.2-8}$$

Certain values of the gamma function can be expressed in K (taken from [259, p.12]):

$$\Gamma\left(\frac{1}{3}\right) \quad = \quad \frac{\pi^{1/3}\,2^{7/9}}{3^{1/12}}\,K\left(\frac{\sqrt{3}-1}{2\sqrt{2}}\right)^{1/3} \tag{30.2-9a}$$

$$\Gamma\left(\frac{1}{4}\right) \quad = \quad \pi^{1/4}\,2\,K\left(\frac{1}{\sqrt{2}}\right)^{1/2} \tag{30.2-9b}$$

$$\Gamma\left(\frac{1}{8}\right) \quad = \quad \pi^{1/8}\,2^{17/8}\,K\left(\frac{1}{\sqrt{2}}\right)^{1/4} K\left(\sqrt{2}-1\right)^{1/2} \tag{30.2-9c}$$

30.2.2 Elliptic E

The function E can be defined as

$$E(k) \quad = \quad \int_0^{\pi/2} \sqrt{1 - k^2\,\sin^2\vartheta}\,d\,\vartheta \;=\; \int_0^1 \frac{\sqrt{1-k^2t^2}}{\sqrt{1-t^2}}\,d\,t \tag{30.2-10}$$

One has

$$E(k) = \frac{\pi}{2}\, F\left(\begin{matrix}-\frac{1}{2},\frac{1}{2}\\ 1\end{matrix}\,\middle|\, k^2\right) \tag{30.2-11a}$$

$$= \frac{\pi}{2}\left(-\sum_{i=0}^{\infty}\left(\frac{(2\,i-1)!!}{2^i\,i!}\right)^2 \frac{k^{2i}}{2i-1}\right) = \frac{\pi}{2}\sum_{i=0}^{\infty}\left(\frac{\binom{2\,i}{i}}{4^i}\right)^2 \frac{k^{2i}}{2\,i-1} \tag{30.2-11b}$$

$$= \frac{\pi}{2}\left(1-\left(\frac{1}{2}\right)^2 k^2 - \left(\frac{1\cdot 3}{2\cdot 4}\right)^2 \frac{k^4}{3} - \left(\frac{1\cdot 3\cdot 5}{2\cdot 4\cdot 6}\right)^2 \frac{k^6}{5} - \ldots\right) \tag{30.2-11c}$$

$$= \frac{\pi}{2}\left(1-\frac{1}{4}\,k^2 - \frac{3}{64}\,k^4 - \frac{5}{256}\,k^6 - \frac{175}{16384}\,k^8 - \frac{441}{65536}\,k^{10} - \ldots\right) \tag{30.2-11d}$$

Special values are $E(0)=\frac{\pi}{2}$ and $E(1)=1$. The latter leads to a (slowly converging) series for $2/\pi$:

$$\frac{2}{\pi} = F\left(\begin{matrix}-\frac{1}{2},\frac{1}{2}\\ 1\end{matrix}\,\middle|\, 1\right) \tag{30.2-12}$$

Similarly as for K', one defines E' as

$$E'(k) := E(k') = E(\sqrt{1-k^2}) \tag{30.2-13}$$

The key to fast computation of E is the relation

$$\frac{E}{K} = 1 - \frac{1}{2}\sum_{n=0}^{\infty} 2^n c'^2_n \tag{30.2-14}$$

The terms c' in the sum occur naturally during the computation of the AGM, see relation 30.1-2 on page 593. One defines

$$R := \frac{E}{K}, \qquad R' := \frac{E'}{K'} \tag{30.2-15}$$

Then E can be computed via

$$E(k) = R(k)\,K(k) = \frac{\pi}{2\ \mathrm{AGM}(1,\sqrt{1-k^2})}\cdot\left(1-\sum_{n=0}^{\infty} 2^{n-1} c'^2_n\right) \tag{30.2-16}$$

Legendre's relation between K and E is (arguments omitted for readability, choose your favorite form):

$$\frac{E}{K}+\frac{E'}{K'}-1 = \frac{\pi}{2\,K\,K'} \tag{30.2-17a}$$

$$E\,K'+E'\,K-K\,K' = \frac{\pi}{2} \tag{30.2-17b}$$

Equivalently,

$$\mathrm{AGM}(1,k) = \frac{E/K}{(1-E'/K')} = \frac{R}{(1-R')} \tag{30.2-18}$$

For $k=\frac{1}{\sqrt{2}} =: s$ we have $k=k'$, thereby $K=K'$ and $E=E'$, so

$$\frac{K(s)}{\pi}\left(\frac{2\,E(s)}{\pi}-\frac{K(s)}{\pi}\right) = \frac{1}{2\,\pi} \tag{30.2-19}$$

As expressions 30.2-3a and 30.2-16 provide a fast AGM based computation of $\frac{K}{\pi}$ and $\frac{E}{\pi}$ the above formula can be used to compute π.

Using $E - K = k\,k'\,\frac{dK}{dk} - k^2\,K$ one can express the derivative of K in terms of E and K and thereby compute that quantity fast:

$$\frac{dK}{dk} = \frac{E - k'^2\,K}{k\,k'^2} \tag{30.2-20}$$

For the derivative of E we have

$$\frac{dE}{dk} = \frac{E - K}{k} \tag{30.2-21}$$

We note the following generalization of Legendre's relation in terms of hypergeometric functions (see section 35.2 on page 685):

$$\begin{aligned}\frac{\Gamma\left(1+a+b\right)\,\Gamma\left(1+c+b\right)}{\Gamma\left(\frac{3}{2}+a+b+c\right)\,\Gamma\left(\frac{1}{2}+b\right)} = \\ +F\left(\begin{matrix}+\frac{1}{2}+a,\ -\frac{1}{2}-c\\ 1+a+b\end{matrix}\middle|\,z\right)\ F\left(\begin{matrix}+\frac{1}{2}-a,\ +\frac{1}{2}+c\\ 1+c+b\end{matrix}\middle|\,1-z\right)+ \\ +F\left(\begin{matrix}+\frac{1}{2}+a,\ +\frac{1}{2}-c\\ 1+a+b\end{matrix}\middle|\,z\right)\ F\left(\begin{matrix}-\frac{1}{2}-a,\ +\frac{1}{2}+c\\ 1+c+b\end{matrix}\middle|\,1-z\right)- \\ -F\left(\begin{matrix}+\frac{1}{2}-a,\ +\frac{1}{2}-c\\ 1+a+b\end{matrix}\middle|\,z\right)\ F\left(\begin{matrix}+\frac{1}{2}-a,\ +\frac{1}{2}+c\\ 1+c+b\end{matrix}\middle|\,1-z\right)\end{aligned} \tag{30.2-22}$$

This equation is given in [15, p.138]. For $a = b = c = 0$ one obtains Legendre's relation.

30.3 AGM-type algorithms for hypergeometric functions

We give AGM based algorithms for $F\left(\begin{matrix}1/2-s,\,1/2+s\\ 1\end{matrix}\middle|\,z\right)$ where $s \in \{0,\,1/6,\,1/4,\,1/3\}$, and $F\left(\begin{matrix}1/4-t,\,1/4+t\\ 1\end{matrix}\middle|\,z\right)$ where $t \in \{1/12,\,1/6\}$. These are taken from [59] and [124], both papers are recommended for further studies. See also [191], [58], [81], and [80]. The limit of a three-term iteration as a generalized hypergeometric function is determined in [172]. A four-term iteration is considered in [57].

The following transformations can be applied to the functions, these are special cases of relations 35.2-33a and 35.2-33b on page 690:

$$F\left(\begin{matrix}\frac{1}{2}+s,\ \frac{1}{2}-s\\ 1\end{matrix}\middle|\,z\right) = F\left(\begin{matrix}\frac{1}{4}+\frac{s}{2},\ \frac{1}{4}-\frac{s}{2}\\ 1\end{matrix}\middle|\,4z(1-z)\right) \quad \text{where} \quad |z| < \frac{1}{2} \tag{30.3-1a}$$

$$F\left(\begin{matrix}\frac{1}{4}+t,\ \frac{1}{4}-t\\ 1\end{matrix}\middle|\,z\right) = F\left(\begin{matrix}\frac{1}{2}+2t,\ \frac{1}{2}-2t\\ 1\end{matrix}\middle|\,\frac{1-\sqrt{1-z}}{2}\right) \tag{30.3-1b}$$

Algorithms for $F\left(\begin{matrix}1/2,\,1/2\\ 1\end{matrix}\middle|\,z\right)$ $[1/2 \pm 0]$

The following is relation 30.2-3a on page 595, the classical AGM algorithm which has quadratic convergence:

$$F\left(\begin{matrix}1/2,\ 1/2\\ 1\end{matrix}\middle|\,z\right) = 1/M\left(1,\,\sqrt{1-z}\right) \quad \text{where} \tag{30.3-2a}$$

$$M(a,b) := \left[(a+b)/2,\,\sqrt{a\,b}\right] \tag{30.3-2b}$$

We write the AGM as $M := [f(a,b),\,g(a,b)]$ in the obvious way.

A fourth order algorithm obtained by combining two steps of the classical AGM:

$$F\left(\begin{matrix}1/2,\,1/2\\1\end{matrix}\Big|\,z\right) \;=\; 1/M\left(1,\,\sqrt[4]{1-z}\right)^2 \quad \text{where} \tag{30.3-3a}$$

$$M(a,b) \;:=\; \left[(a+b)/2,\;\sqrt{a\,b\,(a^2+b^2)/2}\right] \tag{30.3-3b}$$

For comparison, we give the quadratic transform for the hypergeometric function

$$F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\1\end{matrix}\Big|\,z'\right) \;=\; (1+z)\,F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\1\end{matrix}\Big|\,z^2\right) \tag{30.3-4a}$$

where

$$z \;=\; \frac{1-(1-z')^{1/2}}{1+(1-z')^{1/2}} \tag{30.3-4b}$$

$$z' \;=\; 1-\left(\frac{1-z}{1+z}\right)^2 \tag{30.3-4c}$$

It is the special case $a=1/2$ and $b=1/2$ of the transformation

$$F\left(\begin{matrix}a,\,b\\2b\end{matrix}\Big|\,\frac{4z}{(1+z)^2}\right) \;=\; F\left(\begin{matrix}a,\,b\\2b\end{matrix}\Big|\,1-\left(\frac{1-z}{1+z}\right)^2\right) \;=\; (1+z)^{2a}\,F\left(\begin{matrix}a,\;a-b+\frac{1}{2}\\b+\frac{1}{2}\end{matrix}\Big|\,z^2\right) \tag{30.3-5}$$

Algorithms for $F\left(\begin{matrix}1/3,\,2/3\\1\end{matrix}\Big|\,z\right)$ $[1/2 \pm 1/6]$

A third order algorithm:

$$F\left(\begin{matrix}1/3,\,1/3\\1\end{matrix}\Big|\,z\right) \;=\; 1/M\left(\sqrt[3]{1-z},\,1\right) \;=\; \sqrt[3]{1-z}\,F\left(\begin{matrix}2/3,\,2/3\\1\end{matrix}\Big|\,z\right) \quad \text{where} \tag{30.3-6a}$$

$$M(a,b) \;:=\; \left[(a+2\,b)/3,\;\sqrt[3]{b\,(a^2+a\,b+b^2)\,/3}\right] \tag{30.3-6b}$$

One further has

$$F\left(\begin{matrix}1/3,\,2/3\\1\end{matrix}\Big|\,z\right) \;=\; 1/M\left(1,\,\sqrt[3]{1-z}\right) \tag{30.3-7}$$

A quadratic algorithm:

$$F\left(\begin{matrix}1/3,\,2/3\\1\end{matrix}\Big|\,z\right) \;=\; 1/M\left(1,\,\sqrt[3]{1-z}\right) \quad \text{where} \tag{30.3-8a}$$

$$M(a,b) \;:=\; \left[\frac{1}{2}\left(\sqrt[3]{2p-a^3}+\sqrt[3]{2m-a^3}\right),\;\frac{1}{2}\left(\sqrt[3]{p}+\sqrt[3]{m}\right)\right] \qquad \text{and} \tag{30.3-8b}$$

$$p \;:=\; b^3+t, \qquad m \;:=\; b^3-t, \qquad t \;:=\; \sqrt{b^6-a^3\,b^3} \tag{30.3-8c}$$

And again (see relation 30.3-6a):

$$F\left(\begin{matrix}1/3,\,1/3\\1\end{matrix}\Big|\,z\right) \;=\; 1/M\left(\sqrt[3]{1-z},\,1\right) \;=\; \sqrt[3]{1-z}\,F\left(\begin{matrix}2/3,\,2/3\\1\end{matrix}\Big|\,z\right) \tag{30.3-9}$$

We note the following hypergeometric transformation due to Ramanujan:

$$F\left(\begin{matrix}\frac{1}{3},\,\frac{2}{3}\\1\end{matrix}\Big|\,z'\right) \;=\; (1+2\,z)\,F\left(\begin{matrix}\frac{1}{3},\,\frac{2}{3}\\1\end{matrix}\Big|\,z^3\right) \tag{30.3-10a}$$

where

$$z = \frac{1-(1-z')^{1/3}}{1+2\,(1-z')^{1/3}} \tag{30.3-10b}$$

$$z' = 1-\left(\frac{1-z}{1+2z}\right)^3 \tag{30.3-10c}$$

The general form is given in [39]:

$$F\left(\begin{matrix}c,\ c+\frac{1}{3}\\ \frac{3c+1}{2}\end{matrix}\middle|\ 1-\left(\frac{1-z}{1+2z}\right)^3\right) = (1+2z)^{3c}\,F\left(\begin{matrix}c,\ c+\frac{1}{3}\\ \frac{3c+5}{6}\end{matrix}\middle|\ z^3\right) \tag{30.3-11}$$

For $c = 1/3$ one obtains relation 30.3-10a. A computer algebra proof that relation 30.3-11 is the only possible generalization of relation 30.3-10c is given in [171].

An alternative quadratic algorithm is

$$F\left(\begin{matrix}1/3,\ 2/3\\ 1\end{matrix}\middle|\ z\right) = 1/M\,(1,\,W) \quad \text{where} \tag{30.3-12a}$$

$$M(a,b) := \left[(a+b)/2,\ \left(3\,\sqrt{b\,(b+2a)/3}-b\right)/2\right] \quad \text{and} \tag{30.3-12b}$$

$$W := \frac{1-R+R^2}{R},\quad R := \left[\sqrt{u^2-1}+u\right]^{1/3},\quad u := 1-2\,z \tag{30.3-12c}$$

It is given in the form

$$F\left(\begin{matrix}1/3,\ 2/3\\ 1\end{matrix}\middle|\ (1-x)\,(1+x/2)^2\right) = 1/M\,(1,\,x) \tag{30.3-13}$$

A product form can be derived from [38, Theorem 6.1]: Let

$$\alpha(z) := \frac{z\,(3+z)^2}{2\,(1+z)^3} \tag{30.3-14a}$$

$$p(z) := \frac{r^2-r+1}{r} \quad \text{where} \quad r := \left[2z+2\sqrt{z^2-z}-1\right]^{1/3} \tag{30.3-14b}$$

then, with $t_0 := 1-z$ and $t_{k+1} := \alpha(p(t_k))$,

$$F\left(\begin{matrix}1/3,\ 2/3\\ 1\end{matrix}\middle|\ z\right) = \left[\prod_{k=0}^{\infty}\frac{1+p(t_k)}{2}\right]^{-1} \tag{30.3-14c}$$

Convergence is quadratic. The function $p(z)$ is the real solution of $p(\beta(z)) = z$ where $\beta(z) := (z^2(3+z))/4$.

Algorithms for $F\left(\begin{matrix}1/4,\,3/4\\ 1\end{matrix}\middle|\ z\right)$ $[1/2 \pm 1/4]$

A quadratic algorithm:

$$F\left(\begin{matrix}1/4,\ 3/4\\ 1\end{matrix}\middle|\ z\right) = 1/M\left(1,\,\sqrt{1-z}\right)^{1/2} \quad \text{where} \tag{30.3-15a}$$

$$M(a,b) := \left[(a+3b)/4,\ \sqrt{b\,(a+b)/2}\right] \tag{30.3-15b}$$

One further has (note the swapped arguments in the mean)

$$F\left(\begin{matrix}1/4,\ 1/4\\ 1\end{matrix}\middle|\ z\right) = 1/M\left(\sqrt{1-z},\,1\right)^{1/2} = \sqrt{1-z}\,F\left(\begin{matrix}3/4,\ 3/4\\ 1\end{matrix}\middle|\ z\right) \tag{30.3-16}$$

Now set $A_k := \sqrt{(a_k + b_k)/2}$ and $B_k := \sqrt{b_k}$, then

$$A_{k+1} = \frac{1}{2}\,(A_k + B_k) \tag{30.3-17a}$$

$$B_{k+1} = \sqrt{A_k\,B_k} \tag{30.3-17b}$$

This is the iteration of the classical AGM. Thereby

$$M(a,b)^{1/2} = \text{AGM}\left(\sqrt{\frac{a+b}{2}},\,\sqrt{b}\right) \tag{30.3-18}$$

and one can employ the AGM scheme of Schönhage (relations 30.1-4a… 30.1-4k on page 594). Equivalently,

$$F\left(\begin{matrix}1/4,\,3/4\\1\end{matrix}\middle|\,z\right) = 1/\,\text{AGM}\left(\sqrt{\frac{1+\sqrt{1-z}}{2}},\,\sqrt[4]{1-z}\right) \tag{30.3-19}$$

We also have

$$F\left(\begin{matrix}1/4,\,1/4\\1\end{matrix}\middle|\,z\right) = 1/\,\text{AGM}\left(\sqrt{\frac{1+\sqrt{1-z}}{2}},\,1\right) \tag{30.3-20}$$

For comparison, we give the hypergeometric transformation

$$F\left(\begin{matrix}\frac{1}{4},\,\frac{3}{4}\\1\end{matrix}\middle|\,z'\right) = \sqrt{1+3\,z}\;F\left(\begin{matrix}\frac{1}{4},\,\frac{3}{4}\\1\end{matrix}\middle|\,z^2\right) \tag{30.3-21a}$$

where

$$z = \frac{1-(1-z')^{1/2}}{1+3\,(1-z')^{1/2}} \tag{30.3-21b}$$

$$z' = 1-\left(\frac{1-z}{1+3\,z}\right)^2 \tag{30.3-21c}$$

is the special case $d = 1/4$ of the transformation

$$F\left(\begin{matrix}d,\,d+\frac{1}{2}\\\frac{4d+2}{3}\end{matrix}\middle|\,z'\right) = (1+3\,z)^{2d}\;F\left(\begin{matrix}d,\,d+\frac{1}{2}\\\frac{4d+5}{6}\end{matrix}\middle|\,z^2\right) \tag{30.3-22}$$

This is taken from [125] where various such transformations and their generalizations are given.

Algorithm for $F\left(\begin{matrix}1/6,\,1/3\\1\end{matrix}\middle|\,z\right)$ $[1/4 \pm 1/12]$

A quadratic algorithm is

$$F\left(\begin{matrix}1/6,\,1/3\\1\end{matrix}\middle|\,z\right) = 1/M\,(1,\,W)^{1/2} \quad \text{where} \tag{30.3-23a}$$

$$M(a,b) := \Big[(a+3b)/4,\,(\sqrt{ab}+b)/2\Big] \qquad \text{and} \tag{30.3-23b}$$

$$W := \frac{1+R+R^2}{3\,R},\quad R := \Big[\sqrt{z^2-1}+z\Big]^{1/3} \tag{30.3-23c}$$

It is given [59, p.515] as

$$\left[F\left(\begin{matrix}1/6,\,1/3\\1\end{matrix}\middle|\,27\,x^2\,(1-x)/4\right)\right]^2 = 1/M\,(1,\,x) \tag{30.3-24}$$

One can solve for the argument of F to obtain the explicit form.

Algorithm for $F\left(\begin{matrix}1/12,\,5/12\\1\end{matrix}\middle|\,z\right)$ $[1/4 \pm 1/6]$

The following algorithm has quadratic convergence, W is defined by relation 30.3-23c:

$$F\left(\begin{matrix}1/12,\,5/12\\1\end{matrix}\middle|\,z\right) = 1/M\left(1,\,W\right)^{1/4} \quad \text{where} \tag{30.3-25a}$$

$$M(a,b) := \left[(a+15b)/16,\,\left(\sqrt{b\,(a+3b)/4}+b\right)/2\right] \tag{30.3-25b}$$

The next relation is the special case $a = 1/6$ of relation 35.2-42e on page 691:

$$F\left(\begin{matrix}\frac{1}{6},\,\frac{1}{6}\\1\end{matrix}\middle|\,z\right) = (1-z)^{-1/6}\,F\left(\begin{matrix}\frac{1}{12},\,\frac{5}{12}\\1\end{matrix}\middle|\,\frac{-4\,z}{(1-z)^2}\right) \tag{30.3-26}$$

The following relations are given in [191, p.17]:

$$F\left(\begin{matrix}\frac{1}{4},\,\frac{1}{4}\\1\end{matrix}\middle|\,-\frac{z}{64}\right) = \left[\frac{1}{16^3}\,(z+16)^3\right]^{-1/12} F\left(\begin{matrix}\frac{1}{12},\,\frac{5}{12}\\1\end{matrix}\middle|\,\frac{1728\,z}{(z+16)^3}\right) \tag{30.3-27a}$$

$$F\left(\begin{matrix}\frac{1}{3},\,\frac{1}{3}\\1\end{matrix}\middle|\,-\frac{z}{27}\right) = \left[\frac{1}{3^6}\,(z+3)^3\,(z+27)\right]^{-1/12} F\left(\begin{matrix}\frac{1}{12},\,\frac{5}{12}\\1\end{matrix}\middle|\,\frac{1728\,z}{(z+3)^3\,(z+27)}\right) \tag{30.3-27b}$$

$$F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\1\end{matrix}\middle|\,-\frac{z}{16}\right) = \left[\frac{1}{16^3}\,(z^2+16z+16)^3\right]^{-1/12} F\left(\begin{matrix}\frac{1}{12},\,\frac{5}{12}\\1\end{matrix}\middle|\,\frac{1728\,z\,(z+16)}{(z^2+16z+16)^3}\right) \tag{30.3-27c}$$

30.4 Computation of π

We give various iterations for computing π with super-linear convergence. The number of full precision multiplications (FPM) is an indication of the efficiency of the algorithm. The approximate number of FPMs that were counted with a computation of π to 4 million decimal digits (using radix $10,000$ and 1 million LIMBs) is indicated like this: #FPM=123.4. The number is computed as three times the FFT-work in terms of full precision real valued FFTs.

30.4.1 Super-linear iterations for π

AGM implemented in [hfloat: src/pi/piagm.cc], #FPM=98.4:

$$a_0 = 1, \qquad b_0 = \frac{1}{\sqrt{2}} \tag{30.4-1a}$$

$$a_{k+1} = \frac{a_k + b_k}{2} \tag{30.4-1b}$$

$$b_{k+1} = \sqrt{a_k\,b_k} \tag{30.4-1c}$$

$$p_n = \frac{2\,a_{n+1}^2}{1-\sum_{k=0}^{n} 2^k\,c_k^2} \quad \to \pi \tag{30.4-1d}$$

$$\pi - p_n = \frac{\pi^2\,2^{n+4}\,e^{-\pi\,2^{n+1}}}{\mathrm{AGM}^2(a_0,b_0)} \tag{30.4-1e}$$

Convergence is second order. Computing π based on the fourth order AGM (relations 30.1-9a... 30.1-9e on page 594) is possible by setting the second argument of the routine (#FPM=149.3 for the quartic variant). Schönhage's variant of the AGM computation (relations 30.1-4a... 30.1-4k on page 594) is implemented in [hfloat: src/pi/piagmsch.cc] (#FPM=78.424).

The AGM method goes back to Gauss, a facsimile of the entry in his 1809 handbook 6 is given in [17, p.101]. The entry states that

$$\pi \;=\; \frac{\mathrm{AGM}(1,k)\;\mathrm{AGM}(1,k')}{1-\sum_{k=0}^{\infty}2^{k-1}\,(c_k^2+c'^2_k)} \tag{30.4-2}$$

where $k'=b_0/a_0$ and $k=\sqrt{1-b_0^2/a_0^2}=c_0/a_0$. For $k=k'=1/\sqrt{2}$ one obtains relation 30.4-1d. The formula appeared also 1924 in [163, p.39]. The algorithm was rediscovered 1976 independently by Brent [64] (reprinted in [36, p.424]) and Salamin [227] (reprinted in [36, p.418]).

AGM variant given in [54], [hfloat: src/pi/piagm3.cc], #FPM=99.5 (#FPM=155.3 for the quartic variant):

$$a_0 \;=\; 1, \qquad b_0 \;=\; \frac{\sqrt{6}+\sqrt{2}}{4} \tag{30.4-3a}$$

$$p_n \;=\; \frac{2\,a_{n+1}^2}{\sqrt{3}\,(1-\sum_{k=0}^{n}2^k\,c_k^2)-1} \quad\to\pi \tag{30.4-3b}$$

$$\pi-p_n \;<\; \frac{\sqrt{3}\,\pi^2\,2^{n+4}\,e^{-\sqrt{3}\,\pi\,2^{n+1}}}{\mathrm{AGM}^2(a_0,b_0)} \tag{30.4-3c}$$

AGM variant given in [54], [hfloat: src/pi/piagm3.cc], #FPM=108.2 (#FPM=169.5 for the quartic variant):

$$a_0 \;=\; 1, \qquad b_0 \;=\; \frac{\sqrt{6}-\sqrt{2}}{4} \tag{30.4-4a}$$

$$p_n \;=\; \frac{6\,a_{n+1}^2}{\sqrt{3}\,(1-\sum_{k=0}^{n}2^k\,c_k^2)+1} \quad\to\pi \tag{30.4-4b}$$

$$\pi-p_n \;<\; \frac{\frac{1}{\sqrt{3}}\,\pi^2\,2^{n+4}\,e^{-\frac{1}{\sqrt{3}}\,\pi\,2^{n+1}}}{\mathrm{AGM}(a_0,b_0)^2} \tag{30.4-4c}$$

Second order iteration from [56, p.170], [hfloat: src/pi/pi2nd.cc], #FPM=255.7:

$$y_0 \;=\; \frac{1}{\sqrt{2}}, \qquad a_0 \;=\; \frac{1}{2} \tag{30.4-5a}$$

$$y_{k+1} \;=\; \frac{1-(1-y_k^2)^{1/2}}{1+(1-y_k^2)^{1/2}} \quad\to 0+ \tag{30.4-5b}$$

$$\;=\; \frac{(1-y_k^2)^{-1/2}-1}{(1-y_k^2)^{-1/2}+1} \tag{30.4-5c}$$

$$a_{k+1} \;=\; a_k\,(1+y_{k+1})^2-2^{k+1}\,y_{k+1} \quad\to\frac{1}{\pi} \tag{30.4-5d}$$

$$a_k-\pi^{-1} \;\le\; 16\cdot 2^{k+1}\,e^{-2^{k+1}\,\pi} \tag{30.4-5e}$$

Relation 30.4-5c shows how to save 1 multiplication per step (see section 28.1 on page 559). A simple proof of this iteration is given in [139].

Borwein's quartic (fourth order) iteration from [56, p.170], variant $r=4$, implemented in [hfloat:

src/pi/pi4th.cc], #FPM=170.5:

$$\begin{aligned}
y_0 &= \sqrt{2}-1, \qquad a_0 = 6-4\sqrt{2} && (30.4\text{-}6a)\\
y_{k+1} &= \frac{1-(1-y_k^4)^{1/4}}{1+(1-y_k^4)^{1/4}} \quad \to 0+ && (30.4\text{-}6b)\\
&= \frac{(1-y_k^4)^{-1/4}-1}{(1-y_k^4)^{-1/4}+1} && (30.4\text{-}6c)\\
a_{k+1} &= a_k\,(1+y_{k+1})^4 - 2^{2k+3}\,y_{k+1}\,(1+y_{k+1}+y_{k+1}^2) \quad \to \frac{1}{\pi} && (30.4\text{-}6d)\\
&= a_k\,((1+y_{k+1})^2)^2 - 2^{2k+3}\,y_{k+1}\,((1+y_{k+1})^2-y_{k+1}) && (30.4\text{-}6e)\\
0 &< a_k-\pi^{-1} \le 16\cdot 4^n\,2\,e^{-4^n\,2\,\pi} && (30.4\text{-}6f)
\end{aligned}$$

Identities 30.4-6c and 30.4-6e show how to save operations.

Borwein's quartic (fourth order) iteration, variant $r = 16$, implemented in [hfloat: src/pi/pi4th.cc], #FPM=164.4:

$$\begin{aligned}
y_0 &= \frac{1-2^{-1/4}}{1+2^{-1/4}}, \qquad a_0 = \frac{8/\sqrt{2}-2}{(2^{-1/4}+1)^4} && (30.4\text{-}7a)\\
y_{k+1} &= \frac{(1-y_k^4)^{-1/4}-1}{(1-y_k^4)^{-1/4}+1} \quad \to 0+ && (30.4\text{-}7b)\\
a_{k+1} &= a_k\,(1+y_{k+1})^4 - 2^{2k+4}\,y_{k+1}\,(1+y_{k+1}+y_{k+1}^2) \quad \to \frac{1}{\pi} && (30.4\text{-}7c)\\
0 &< a_k-\pi^{-1} \le 16\cdot 4^n\,4\,e^{-4^n\,4\,\pi} && (30.4\text{-}7d)
\end{aligned}$$

The operation count is unchanged, but this variant gives approximately twice as much precision after the same number of steps. The general form of the quartic iterations (relations 30.4-6a..., and 30.4-7a...) is given in [56, pp.170ff]:

$$\begin{aligned}
y_0 &= \sqrt{\lambda^*(r)}, \qquad a_0 = \alpha(r) && (30.4\text{-}8a)\\
y_{k+1} &= \frac{(1-y_k^4)^{-1/4}-1}{(1-y_k^4)^{-1/4}+1} \quad \to 0+ && (30.4\text{-}8b)\\
a_{k+1} &= a_k\,(1+y_{k+1})^4 - 2^{2k+2}\,\sqrt{r}\,y_k\,(1+y_{k+1}+y_{k+1}^2) \quad \to \frac{1}{\pi} && (30.4\text{-}8c)\\
0 &< a_k-\pi^{-1} \le \quad 16\cdot 4^n\,\sqrt{r}\,e^{-4^n\,\sqrt{r}\,\pi} && (30.4\text{-}8d)
\end{aligned}$$

Derived AGM iteration (second order, from [56, pp.46ff]), implemented in [hfloat: src/pi/pideriv.cc], #FPM=276.2:

$$\begin{aligned}
x_0 &= \sqrt{2}, \qquad p_0 = 2+\sqrt{2}, \qquad y_1 = 2^{1/4} && (30.4\text{-}9a)\\
x_{k+1} &= \frac{1}{2}\left(\sqrt{x_k}+\frac{1}{\sqrt{x_k}}\right) \quad (k\ge 0) \quad \to 1+ && (30.4\text{-}9b)\\
y_{k+1} &= \frac{y_k\,\sqrt{x_k}+\dfrac{1}{\sqrt{x_k}}}{y_k+1} \quad (k\ge 1) \quad \to 1+ && (30.4\text{-}9c)\\
p_{k+1} &= p_k\,\frac{x_k+1}{y_k+1} \quad (k\ge 1) \quad \to \pi+ && (30.4\text{-}9d)\\
p_k-\pi &< 10^{-2^{k+1}} && (30.4\text{-}9e)
\end{aligned}$$

Cubic AGM iteration (third order) from [60], implemented in [hfloat: src/pi/picubagm.cc], #FPM=182.7:

$$a_0 = 1, \qquad b_0 = \frac{\sqrt{3}-1}{2} \tag{30.4-10a}$$

$$a_{n+1} = \frac{a_n + 2\,b_n}{3} \tag{30.4-10b}$$

$$b_{n+1} = \sqrt[3]{\frac{b_n\,(a_n^2 + a_n\,b_n + b_n^2)}{3}} \tag{30.4-10c}$$

$$p_n = \frac{3\,a_n^2}{1 - \sum_{k=0}^{n} 3^k\,(a_k^2 - a_{k+1}^2)} \quad \to \pi \tag{30.4-10d}$$

Quintic (5th order) iteration from [56, p.310], [hfloat: src/pi/pi5th.cc], #FPM=353.2:

$$s_0 = 5(\sqrt{5}-2), \qquad a_0 = \frac{1}{2} \tag{30.4-11a}$$

$$x = \frac{5}{s_n} - 1 \quad \to 4 \tag{30.4-11b}$$

$$y = (x-1)^2 + 7 \quad \to 16 \tag{30.4-11c}$$

$$z = \left(\frac{x}{2}\,\left(y + \sqrt{y^2 - 4x^3}\right)\right)^{1/5} \quad \to 2 \tag{30.4-11d}$$

$$s_{n+1} = \frac{25}{s_n(z + x/z + 1)^2} \quad \to 1 \tag{30.4-11e}$$

$$a_{n+1} = s_n^2 a_n - 5^n \left(\frac{s_n^2 - 5}{2} + \sqrt{s_n\,(s_n^2 - 2s_n + 5)}\right) \quad \to \frac{1}{\pi} \tag{30.4-11f}$$

$$a_n - \frac{1}{\pi} < 16 \cdot 5^n\, e^{-\pi\,5^n} \tag{30.4-11g}$$

Cubic (third order) iteration from [25], implemented in [hfloat: src/pi/pi3rd.cc], #FPM=200.3:

$$a_0 = \frac{1}{3}, \qquad s_0 = \frac{\sqrt{3}-1}{2} \tag{30.4-12a}$$

$$r_{k+1} = \frac{3}{1 + 2\,(1 - s_k^3)^{1/3}} \tag{30.4-12b}$$

$$s_{k+1} = \frac{r_{k+1} - 1}{2} \tag{30.4-12c}$$

$$a_{k+1} = r_{k+1}^2\,a_k - 3^k\,(r_{k+1}^2 - 1) \quad \to \frac{1}{\pi} \tag{30.4-12d}$$

Nonic (9th order) iteration from [25], implemented in [hfloat: src/pi/pi9th.cc], #FPM=273.7:

$$a_0 = \frac{1}{3}, \qquad r_0 = \frac{\sqrt{3}-1}{2}, \qquad s_0 = (1 - r_0^3)^{1/3} \tag{30.4-13a}$$

$$t = 1 + 2\,r_k \tag{30.4-13b}$$

$$u = \left(9\,r_k\,(1 + r_k + r_k^2)\right)^{1/3} \tag{30.4-13c}$$

$$v = t^2 + t\,u + u^2 \tag{30.4-13d}$$

$$m = \frac{27\,(1 + s_k + s_k^2)}{v} \tag{30.4-13e}$$

$$a_{k+1} = m\,a_k + 3^{2\,k-1}\,(1 - m) \quad \to \frac{1}{\pi} \tag{30.4-13f}$$

$$s_{k+1} = \frac{(1 - r_k)^3}{(t + 2\,u)\,v} \tag{30.4-13g}$$

$$r_{k+1} = (1 - s_k^3)^{1/3} \tag{30.4-13h}$$

30.4.2 Measured timings and operation counts

#FPM	order	routine in hfloat	time
78.424	2	pi_agm_sch()	76 sec
98.424	2	pi_agm()	93 sec
99.510	2	pi_agm3(fast variant)	94 sec
108.241	2	pi_agm3(slow variant)	103 sec
149.324	4	pi_agm(quartic)	139 sec
155.265	4	pi_agm3(quartic, fast variant)	145 sec
164.359	4	pi_4th_order(r=16 variant)	154 sec
169.544	4	pi_agm3(quartic, slow variant)	159 sec
170.519	4	pi_4th_order(r=4 variant)	160 sec
182.710	3	pi_cubic_agm()	173 sec
200.261	3	pi_3rd_order()	189 sec
255.699	2	pi_2nd_order()	240 sec
273.763	9	pi_9th_order()	256 sec
276.221	2	pi_derived_agm()	259 sec
353.202	5	pi_5th_order()	329 sec

Figure 30.4-A: Measured operations counts and timings for various iterations for the computation of π to 4 million decimal digits.

The operation counts and timings for the algorithms given so far when computing π to 4 million decimal digits (using 1 million `LIMB`s and radix 10,000) are shown in figure 30.4-A. In view of these figures it seems surprising that the quartic algorithms `pi_4th_order()` and the quartic AGM `pi_agm(quartic)` are usually considered close competitors to the second order AGM schemes.

Apart from the operation count the number of variables used has to be taken into account. The algorithms using more variables (like `pi_5th_order()`) cannot be used to compute as many digits as those using only a few (notably the AGM-schemes) given a fixed amount of RAM. Higher order algorithms tend to require more variables.

A further disadvantage of the algorithms of higher order is the more discontinuous growth of the work: if just a few more digits are to be computed than are available after step k then an additional step is required. Consider an extreme case where an algorithm T of order $1,000$ would compute 1 million digits after the second step, at a slightly lower cost than the most effective competitor. Then algorithm T would likely be the 'best' one only for small ranges in the number of digits around the values 10^3, 10^6, 10^9, ...

Finally, it is much easier to find special arithmetical optimizations for the 'simple' (low order) algorithms, Schönhage's AGM variant being the prime example.

30.4.3 More iterations for π

The following iterations are not implemented in hfloat.

Third order algorithm from [53]:

$$v_0 = 2^{-1/8}, \qquad v_1 = 2^{-7/8}\left((1-3^{1/2})\,2^{-1/2}+3^{1/4}\right) \tag{30.4-14a}$$

$$w_0 = 1, \qquad \alpha_0 = 1, \qquad \beta_0 = 0 \tag{30.4-14b}$$

$$v_{n+1} = v_n^3 - \left\{v_n^6 + \left[4v_n^2(1-v_n^8)\right]^{1/3}\right\}^{1/2} + v_{n-1} \tag{30.4-14c}$$

$$w_{n+1} = \frac{2v_n^3 + v_{n+1}\left(3v_{n+1}^2v_n^2-1\right)}{2v_{n+1}^3 - v_n\left(3v_{n+1}^2v_n^2-1\right)}\,w_n \tag{30.4-14d}$$

$$\alpha_{n+1} = \left(\frac{2v_{n+1}^3}{v_n}+1\right)\alpha_n \tag{30.4-14e}$$

$$\beta_{n+1} = \left(\frac{2v_{n+1}^3}{v_n}+1\right)\beta_n + \left(6w_{n+1}v_n - 2v_{n+1}w_n\right)\frac{v_{n+1}^2\alpha_n}{v_n^2} \tag{30.4-14f}$$

$$\pi_n = \frac{8\cdot 2^{1/8}}{\alpha_n\,\beta_n} \quad \to \pi \tag{30.4-14g}$$

Second order algorithm from [61]:

$$\alpha_0 = 1/3, \qquad m_0 = 2 \tag{30.4-15a}$$

$$m_{n+1} = \frac{4}{1+\sqrt{(4-m_n)\,(2+m_n)}} \tag{30.4-15b}$$

$$\alpha_{n+1} = m_n\,\alpha_n + \frac{2^n}{3}(1-m_n) \quad \to \frac{1}{\pi} \tag{30.4-15c}$$

Implicit second order algorithm from [61] (also in [58, p.700]):

$$\alpha_0 = 1/3, \qquad s_1 = 1/3 \tag{30.4-16a}$$

$$(s_n)^2 + (s_n^*)^2 = 1 \tag{30.4-16b}$$

$$(1+3\,s_{n+1})\,(1+3\,s_n^*) = 4 \tag{30.4-16c}$$

$$\alpha_{n+1} = (1+3\,s_{n+1})\,\alpha_n - 2^n\,s_{n+1} \quad \to \frac{1}{\pi} \tag{30.4-16d}$$

It is trivial to turn this algorithm into an explicit form as with the next algorithm. However, there exist iterations that cannot be turned into explicit forms.

Implicit fourth order algorithm from [61] (also in [58, p.700]):

$$\alpha_0 = 1/3 \qquad s_1 = \sqrt{2}-1 \tag{30.4-17a}$$

$$(s_n)^4 + (s_n^*)^4 = 1 \tag{30.4-17b}$$

$$(1+3\,s_{n+1})\,(1+3\,s_n^*) = 2 \tag{30.4-17c}$$

$$\alpha_{n+1} = (1+s_{n+1})^4\alpha_n + \frac{4^{n+1}}{3}\left[1-(1+s_{n+1})^4\right] \quad \to \frac{1}{\pi} \tag{30.4-17d}$$

Combining two steps of the fourth order iteration leads to an algorithm of order 16. The following form

is given in [61, p.111]:

$$\alpha_0 = 1/3, \qquad s_1 = \sqrt{2}-1 \tag{30.4-18a}$$
$$s_n^* = \left(1-s_n^4\right)^{1/4} \tag{30.4-18b}$$
$$x_n = 1/\left(1+s_n^*\right)^4 \tag{30.4-18c}$$
$$y_n = x_n\left(1+s_n\right)^4 \tag{30.4-18d}$$
$$\alpha_n = 16\,y_n\,\alpha_{n-1} + \frac{4^{2n-1}}{3}\left[1-12\,x_n-4\,y_n\right] \quad \to \frac{1}{\pi} \tag{30.4-18e}$$
$$t_n = 1+s_n^* \tag{30.4-18f}$$
$$u_n = \left[8\,s_n^*\left(1+{s_n^*}^2\right)\right]^{1/4} \tag{30.4-18g}$$
$$s_{n+1} = \frac{\left(1-s_n^*\right)^4}{(t+u)^2\left(t^2+u^2\right)} \tag{30.4-18h}$$

Quadratic iteration by Christian Hoffmann, given in [147, p.5]:

$$a_0 = \sqrt{2}, \qquad b_0 = 0, \qquad p_0 = 2+\sqrt{2} \tag{30.4-19a}$$
$$a_{n+1} = \frac{1}{2}\left(\sqrt{a_n}+1/\sqrt{a_n}\right) \quad \to 1+ \tag{30.4-19b}$$
$$b_{n+1} = \sqrt{a_n}\,\frac{b_n+1}{b_n+a_n} \quad \to 1- \tag{30.4-19c}$$
$$p_{n+1} = p_n\,b_{n+1}\,\frac{1+a_{n+1}}{1+b_{n+1}} \quad \to \pi \tag{30.4-19d}$$

Note that relation 30.4-19b deviates from the one given in the cited paper which seems to be incorrect. This is a variant of the iteration given as relations 30.4-9a... 30.4-9e on page 604. The values p_k are identical in both iterations.

Cubic iteration given in [55, p.125]:

$$s_0 = \sqrt{3+2\sqrt{3}}, \qquad a_0 = 1/2 \tag{30.4-20a}$$
$$m_n = 3/s_n \tag{30.4-20b}$$
$$a_{n+1} = \left[\left(s_n^2-1\right)^{1/3}+2\right]^2/s_n \tag{30.4-20c}$$
$$a_{n+1} = m_n^2\,a_n - 3^n\left(m_n^2+2m_n-3\right)/2 \quad \to \frac{1}{\pi} \tag{30.4-20d}$$

The cited paper actually gives a more general form, here we take $N=1$ for simplicity.

Cubic iteration given in [81, p.1506, it-1.2]:

$$t_0 = 1/3, \qquad s_0 = \left(\sqrt{3}-1\right)/2 \tag{30.4-21a}$$
$$s_n = \frac{1-\left(1-s_{n-1}^3\right)^{1/3}}{1+2\left(1-s_{n-1}^3\right)^{1/3}} \tag{30.4-21b}$$
$$= \frac{\left(1-s_{n-1}^3\right)^{-1/3}-1}{\left(1-s_{n-1}^3\right)^{-1/3}+2} \tag{30.4-21c}$$
$$t_n = \left(1+2s_n\right)^2\,t_{n-1} - 3^{n-1}\left(\left(1+2s_n\right)^2-1\right) \quad \to \frac{1}{\pi} \tag{30.4-21d}$$

Note the corrected denominator in relation 30.4-21b (exponent of s_{n-1} is wrongly given as 2).

Quadratic iteration given in [81, p.1507, it-1.3]:

$$k_0 = 0, \qquad s_0 = 1/\sqrt{2} \tag{30.4-22a}$$

$$s_n = \frac{1-\sqrt{1-s_{n-1}^2}}{1+\sqrt{1-s_{n-1}^2}} \tag{30.4-22b}$$

$$k_n = (1+s_n)^2\, k_{n-1} + 2^n\,(1-s_n)\,s_n \quad \to \frac{1}{\pi} \tag{30.4-22c}$$

Cubic iteration given in [81, p.1507, it-1.4]:

$$k_0 = 0, \qquad s_0 = 1/\sqrt[3]{2} \tag{30.4-23a}$$

$$s_n = \frac{1-\sqrt[3]{1-s_{n-1}^3}}{1+2\,\sqrt[3]{1-s_{n-1}^3}} \tag{30.4-23b}$$

$$k_n = (1+2\,s_n)^2\, k_{n-1} + 8\cdot 3^{n-2}\,\sqrt{3}\,s_n\,\frac{1-s_n^3}{1+2\,s_n} \quad \to \frac{1}{\pi} \tag{30.4-23c}$$

Quadratic iteration given in [81, p.1508, it-1.5]:

$$k_0 = 0, \qquad y_0 = 8/9 \tag{30.4-24a}$$

$$y_n = 2\,\frac{6\,y_{n-1}^2 - 5\,y_{n-1} + \sqrt{y_{n-1}(4-y_{n-1})}}{9\,y_{n-1}^2 - 6\,y_{n-1} + 1} \tag{30.4-24b}$$

$$k_n = 2^n\,\sqrt{3}\,\frac{y_{n-1}\,(1-y_{n-1})}{\sqrt{4-3\,y_{n-1}}} + (4-3\,y_{n-1})\,k_{n-1} \quad \to \frac{1}{\pi} \tag{30.4-24c}$$

Quadratic iteration given in [81, p.1508, it-1.6]:

$$k_0 = 0, \qquad y_0 = 4/5 \tag{30.4-25a}$$

$$y_n = \frac{2\,y_{n-1}^2 - y_{n-1} + \sqrt{4\,y_{n-1} - 3\,y_{n-1}^2}}{1+y_{n-1}^2} \tag{30.4-25b}$$

$$Q_n = \sqrt{(y_{n-1}+1)\,(4-3\,y_{n-1})\,(y_{n-1}^2 - 3\,y_{n-1} + 4)} \tag{30.4-25c}$$

$$k_n = \frac{2^n}{\sqrt{7}}\,\frac{y_{n-1}\,(1-y_{n-1})}{2-y_{n-1}}\,Q_n + (2-y_{n-1})^2\,k_{n-1} \quad \to \frac{1}{\pi} \tag{30.4-25d}$$

Quadratic iteration (as product, two forms) given in [57, p.324]:

$$x_1 = \frac{2}{9}\left(\sqrt{6}+2\right), \qquad y_1 = \frac{1}{6}\left(\sqrt{6}+4\right) \tag{30.4-26a}$$

$$x_{n+1} = \frac{2\left(\sqrt{x_n}+x_n\right)}{1+3\,x_n} \tag{30.4-26b}$$

$$y_{n+1} = \frac{2y_n + y_n/\sqrt{x_n} + \sqrt{x_n}}{1+3\,y_n} \tag{30.4-26c}$$

$$\pi = \frac{27}{8}\,\prod_{n=1}^{\infty}\frac{(1+3\,x_n)^2}{(1+3\,y_n)/4} \tag{30.4-26d}$$

$$\pi = \frac{5+2\sqrt{6}}{3}\,\prod_{n=1}^{\infty}\frac{\left(1+1/\sqrt{x_n}\right)^2}{1+3\,y_n} \tag{30.4-26e}$$

The definitive source for iterations to compute π and the underlying mathematics is [56].

30.5 Arctangent relations for π *

This section is about relations of the form

$$k\,\frac{\pi}{4} \quad = \quad m_1 \arctan\frac{1}{x_1} + m_2 \arctan\frac{1}{x_2} + \ldots + m_n \arctan\frac{1}{x_n} \tag{30.5-1}$$

where $k,\ m_1,\ldots,m_n,\ x_1,\ldots,x_n \in \mathbb{Z}$ (in fact, $k = 1$ almost always). This is an n-term relation. For example, a 4-term relation, found 1896 by Størmer [248], is

$$\frac{\pi}{4} \quad = \quad +44 \arctan\frac{1}{57} + 7 \arctan\frac{1}{239} - 12 \arctan\frac{1}{682} + 24 \arctan\frac{1}{12943} \tag{30.5-2}$$

We use the following compact notation

```
m1[x1] +m2[x2] + ... +mn[xn]  == k * Pi/4
```

for relation 30.5-1. For example, Størmer's relation 30.5-2 would be written as

```
+44[57] +7[239] -12[682] +24[12943]   == 1 * Pi/4
```

We write the relations so that the arguments x_j are strictly increasing. Further, n-term relations are sorted so that the first arguments x_1 are in decreasing order (if $x_1 \ldots x_j$ coincide with two relations then the arguments x_{j+1} are used for sorting). For example, a few 6-term relations are

```
+322[577]  +76[682] +139[1393] +156[12943] +132[32807] +44[1049433]   == 1 * Pi/4
+122[319]  +61[378]  +115[557]   +29[1068]   +22[3458]   +44[27493]   == 1 * Pi/4
+100[319] +127[378]   +71[557]   -15[1068]   +66[2943]  +44[478707]   == 1 * Pi/4
+337[307] -193[463] +151[4193]  +305[4246] -122[39307]  -83[390112]   == 1 * Pi/4
+183[268]  +32[682]  +95[1568]   +44[4662] -166[12943]   -51[32807]   == 1 * Pi/4
```

Note that the second and third relation are sorted according to their fifth arguments (3458 and 2943). Among all n-term relations we consider a relation *better* than another if it precedes it. The first one is the *best* relation. Our goal is to find the best n-term relation for n small. For example, the relation

```
+322[577]  +76[682] +139[1393] +156[12943] +132[32807] +44[1049433]   == 1 * Pi/4
```

is the best (known!) 6-term relation. The best n-term relations for $2 \le n \le 12$ currently known are shown in figure 30.5-A. Note that $k = -1$ in the 10-term relation, and $k = 2$ in the 12-term relation. The best relations for $13 \le n \le 21$ (shortened to save space) are shown in figure 30.5-B. Figure 30.5-C gives just the first argument (x_1) of the best relations for $2 \le n \le 27$.

```
+4[5] -1[239]   == 1 * Pi/4

+12[18] +8[57] -5[239]   == 1 * Pi/4

+44[57] +7[239] -12[682] +24[12943]   == 1 * Pi/4

+88[192] +39[239] +100[515] -32[1068] -56[173932]   == 1 * Pi/4

+322[577] +76[682] +139[1393] +156[12943] +132[32807] +44[1049433]   == 1 * Pi/4

+1587[2852] +295[4193] +593[4246] +359[39307]
 +481[55603] +625[211050] -708[390112]   == 1 * Pi/4

+2192[5357] +2097[5507] -227[9466] +832[12943]
 +537[34522] -2287[39307] -171[106007] -708[1115618]   == 1 * Pi/4

+3286[34208] +9852[39307] +5280[41688] +7794[44179]
 +7608[60443] +4357[275807] -1484[390112] -1882[619858] +776[976283]   == 1 * Pi/4

+1106[54193] -30569[78629] -28687[88733] -13882[173932]
 +9127[390112] -9852[478707] -24840[1131527] +4357[3014557]
 +21852[5982670] +23407[201229582]   == -1 * Pi/4

+36462[390112] +135908[485298] +274509[683982] -39581[1984933]
 +178477[2478328] -114569[3449051] -146571[18975991] +61914[22709274]
 -69044[24208144] -89431[201229582] -43938[2189376182]   == 1 * Pi/4

+893758[1049433] +655711[1264557] +310971[1706203] +503625[1984933]
 -192064[2478328] -229138[3449051] -875929[18975991] -616556[21638297]
 -187143[22709274] -171857[24208144] -251786[201229582] -432616[2189376182]   == 2 * Pi/4
```

Figure 30.5-A: Best n-term arctan relations currently known for $2 \le n \le 12$.

```
+1126917[3449051] +1337518[4417548] ... -216308[2189376182]   == 1 * Pi/4
+446879[6826318] +5624457[8082212] ... +483341[17249711432]   == 1 * Pi/4
+5034126[20942043] +1546003[22709274] ... +1337518[250645741818]   == 1 * Pi/4
+14215326[53141564] +6973645[54610269] ... +8735690[34840696582]   == 1 * Pi/4
+12872838[201229582] +27205340[203420807] ... +35839320[134520516108]   == 1 * Pi/4
+2859494[299252491] -41068896[321390012] ... -89623108[18004873694818]   == -1 * Pi/4
+270619381[778401733] -138919506[1012047353] ... +146407224[30038155625330]   == 1 * Pi/4
+807092487[2674664693] +479094776[2701984943] ... +214188292[564340076432]   == 1 * Pi/4
+598245178[5513160193] -115804626[7622130953] ... -1521437626[38057255532937]   == 1 * Pi/4
```

Figure 30.5-B: The best n-term arctan relations (shortened) currently known for $13 \leq n \leq 21$.

```
 n-terms       min-arg
    2                5  Machin (1706)    +4[5] -1[239]   == 1 * Pi/4
    3               18  Gauss (YY?)      +12[18] +8[57] -5[239]   == 1 * Pi/4
    4               57  Stormer (1896)
    5              192  JJ (1993), prev: Stormer (1896) 172
    6              577  JJ (1993)
    7            2,852  JJ (1993)
    8            5,357  JJ (2006), prev: JJ (1993) 4,246
    9           34,208  JJ (2006), prev: JJ (1993) 12,943, prev: Gauss (Y?) 5,257
   10           54,193  JJ (2006), prev: JJ (1993) 51,387
   11          390,112  JJ (1993)
   12        1,049,433  JJ (2006), prev: JJ (1993) 683,982
   13        3,449,051  JJ (2006), prev: JJ (1993) 1,984,933
   14        6,826,318  JJ (2006)
   15       20,942,043  HCL (1997), prev: MRW (1997) 18.975,991
   16       53,141,564  JJ (2006)
   17      201,229,582  JJ (2006)
   18      299,252,491  JJ (2006)
   19      778,401,733  JJ (2006)
   20    2,674,664,693  JJ (2006)
   21    5,513,160,193  JJ (2006)
   22   17,249,711,432  JJ (2006), prev: 16,077,395,443 MRW (27-Jan-2003)
   23   58,482,499,557  JJ (2006)
   24  102,416,588,812  JJ (2006)
   25  160,422,360,532  JJ (2006)
   26  392,943,720,343  JJ (2006)
   27  970,522,492,753  JJ (2006)

MRW := Michael Roby Wetherfield
HCL := Hwang Chien-Lih
JJ := Joerg Arndt
```

Figure 30.5-C: First arguments of the best n-term arctan relation known today, for $2 \leq n \leq 27$. The 4-term entry corresponds to relation 30.5-2.

30.5.1 How to find one relation

In the 5-term relation

```
+88[192] +39[239] +100[515] -32[1068] -56[173932]   == 1 * Pi/4
```

factor $x_j^2 + 1$ for all (inverse) arguments x_j:

```
192^2+1 == 36865 == 5 73 101
239^2+1 == 57122 == 2 13 13 13 13
515^2+1 == 265226 == 2 13 101 101
1068^2+1 == 1140625 == 5 5 5 5 5 73
173932^2+1 == 30252340625 == 5 5 5 5 5 13 73 101 101
```

Note that all odd prime factors are the four primes $5, 13, 73, 101$. The coefficients m_j can be computed as follows. Write (for all arguments x_j)

$$x_j^2 + 1 \quad = \quad 2^{e(j,0)}\, 5^{e(j,1)}\, 13^{e(j,2)}\, 73^{e(j,3)}\, 101^{e(j,4)} \tag{30.5-3}$$

Now define a matrix M using the exponents $e(j, u)$ (ignoring the prime 2):

$$M^T_{j,i} \quad := \quad \pm e(j, i) \tag{30.5-4}$$

The sign of $M_{j,i}$ is minus if $(x_j \bmod p_i) < p_i/2$. With our example we obtain

```
transpose(M) :=
  [-5, -1, +1, -2]  //  173932^2+1 ==   5^5 *13^1 *73^1 *101^2
  [+6,  0, +1,  0]  //  1068^2+1   ==   5^6       *73^1
  [ 0, +1,  0, -2]  //  515^2+1    ==        13^1       *101^2  (*2)
  [ 0, -4,  0,  0]  //  239^2+1    ==        13^4               (*2)
  [-1,  0, +1, +1]  //  192^2+1    ==   5^1       *73^1 *101^1
//  5, 13, 73, 101  <--= primes
```

For the signs of the upper left 3×2 sub-matrix, note that $(173932 \bmod 5) = 2 < 5/2$, $(173932 \bmod 13) = 5 < 13/2$, $(1068 \bmod 5) = 3 > 5/2$, and $(515 \bmod 13) = 8 > 13/2$. The nullspace of M consists of one vector:

```
[-56,   -32,   100,   39,   88]
```

This tells us that

```
+88[192] +39[239] +100[515] -32[1068] -56[173932]   == k * Pi/4
```

We determine that $k = 1$ by a floating point computation of the left hand side. Quite often one finds a relation where $k = 0$, which we are not interested in. For example, the candidates 12943, 1068, 682, 538, 239 all factor into (2 and) the odd primes 5, 13, 61, 73. The matrix M is

```
transpose(M) =
  [+4, +3, -1,  0] //  12943^2+1  ==   5^4 *13^3 *61^1         (*2)
  [+6,  0,  0, +1] //  1068^2+1   ==   5^6             *73^1
  [-3,  0, -2,  0] //  682^2+1    ==   5^3       *61^2
  [+1, -1, +1, -1] //  538^2+1    ==   5^1 *13^1 *61^1 *73^1
  [ 0, -4,  0,  0] //  239^2+1    ==        13^4               (*2)
//  5, 13, 61, 73  <--= primes
```

The nullspace of M is

```
[1, -1, -1, -1, 1]
```

and the relation is

```
+1[239] -1[538] -1[682] -1[1068] +1[12943]   ==   0
```

30.5.2 Searching for sets of candidate arguments

A set of candidate arguments x_j will give a relation only if $x_j^2 + 1$ factor into a common set of primes. Apart from the factor 2, all prime factors are of the form $4\,i+1$. One can choose a subset of those primes $S := \{p_1, \ldots, p_u\}$ and test which of the products $P = 2^{e_0} \cdot p_1^{e_1} \cdots p_u^{e_u}$ are of the form $P = x^2 + 1$. The test is to determine whether $P - 1$ is a perfect square. The pari/gp function `issquare()` does this in an efficient way (as described in [90]). A recursive implementation of the search is

```
1   \\ global variables:
2   ct=0;  \\ count solutions
3   av=vector(1000);  \\ vector containing solutions
4   \\ pv = [...]; \\ vector of primes of the form 4*i+1
5   m=10^20;  \\ search max := sqrt(m)
6
7   check(t)=
8   {
9       local(a);
10      if ( issquare(t-1, &a),  ct++; av[ct] = a; );
11      if ( issquare(t+t-1, &a),  ct++; av[ct] = a; );
12  }
13
14  gen_rec(d, p)=
15  {
16      local(g, gg, t);
17      if ( d>length(pv),  return() );
18      g = pv[d];
19      gg = 1;
20      while ( 1,
21          t = p * gg;
22          if ( t>m,  return() );
23          if ( gg!=1, check(t) );
24          gen_rec(d+1, t);
25          gg *= g;
26      );
27      return();
28  }
```

We do the search using the four primes 5, 13, 61, and 73:

```
pv=[5, 13, 61, 73]; \\ vector of primes
gen_rec(1, 1);  \\ do the search
```

The candidates found are

```
12943, 1068, 682, 538, 239, 57, 27, 18, 11, 8, 7, 5, 3, 2
```

The following relations are found:

```
+1[239] -1[538] -1[682] -1[1068] +1[12943]   ==   NULL
+44[57] +7[538] -5[682] +7[1068] +17[12943]   == 1 * Pi/4  (5-term)
+1[27] +42[57] +6[538] -5[682] +7[1068] +16[12943]   == 1 * Pi/4  (6-term)
+1[18] +41[57] +6[538] -5[682] +6[1068] +16[12943]   == 1 * Pi/4  (6-term)
+1[11] +39[57] +6[538] -5[682] +6[1068] +15[12943]   == 1 * Pi/4  (6-term)
  [--snip--]
+1[3] +26[57] +4[538] -3[682] +4[1068] +10[12943]   == 1 * Pi/4  (6-term)
+1[2] +18[57] +3[538] -2[682] +3[1068] +7[12943]   == 1 * Pi/4  (6-term)
```

The search is reasonably fast for up to about 12 primes. The drawback of this method is that one has to guess which prime set may lead to a good arctan relation. The prime set $\{5, 13, 17, 29, 37, 53, 61, 89, 97, 101\}$ lead (April-1993) to

```
+36462[390112] +135908[485298] +274509[683982] -39581[1984933]
 +178477[2478328] -114569[3449051] -146571[18975991] +61914[22709274]
 -69044[24208144] -89431[201229582] -43938[2189376182]   == 1 * Pi/4
```

which is still the best 11-term relation known today.

The April-2006 computations were done with a more exhaustive search described in the next section.

30.5.3 Exhaustive search for sets of candidate arguments

We want to find all x where $x^2 + 1$ factors into (2 and) the first 64 primes of the form $4\,i + 1$ ($S = \{5, 13, 17, 29, \ldots, 761\}$). Call the resulting set of candidates A. We will later try (for small n) all $(n-1)$-subsets of S and test whether the corresponding subset of A leads to an arctan relation.

The most simple approach is to factor (for x up to a practical maximum) all $x^2 + 1$ and add x to the set A if all odd prime factors of $x^2 + 1$ are in S. The method, however, is rather slow: about 11,000 CPU cycles are needed for each test.

A much faster approach is the following sieving method. We can determine exactly when a prime p divides x^2+1 by solving $x^2 \equiv -1 \pmod{p}$ as shown in section 37.9 on page 773. We can further solve $x^2 \equiv -1 \pmod{p^h}$ for all h as shown in the cited section. Initialize an array with the value 1 for even indices, else with 2 (x^2+1 is even exactly if x is odd). For each prime $p \in S$ do, for all powers p^h, as follows: multiply the array entries with indices $s, s+p^h, s+2\,p^h, s+3\,p^h, \ldots$ where $s^2 \equiv -1 \pmod{p^h}$ by p. Finally find the entries with index x that are equal to x^2+1, these are the candidates.

There are three useful improvements. Firstly, we can use the logarithm of a prime and add it instead of multiplying by the prime. The final test is then whether entry x is (approximately) equal to $\log(x^2+1)$. Secondly, the array can be avoided altogether by using priority queues (see section 4.5 on page 157). An event of the priority queue will be to add $\log(p)$ (initially at index $x = s$ where $s^2 \equiv -1 \pmod{p^h}$). The event must then be rescheduled to $x + p^h$. Thirdly, almost all computations of the logarithm can be avoided by observing that both x^2+1 and the logarithm are strictly increasing functions. We call a number x so that x^2+1 has all odd prime factors in S is a candidate. The sum of logarithms (of primes) for candidates x are equal to $\log(x^2+1)$. If a was the last candidate then for the next candidate b the sum of logarithms must be strictly greater than that for a. Thereby we only need to compute $\log(x^2+1)$ if a new sum of logarithms is greater than the one for the candidate found most recently. It turns out that a logarithm is computed *exactly* whenever a new candidate is found.

The search costs about 250 cycles per test, which is a good improvement over the first attempt. Analysis of the machine code shows that most of the time is spent in the reschedule operations.

The final improvement comes from the separation of the frequent event (small prime powers) from the rare events (big prime powers). One has to use an array again (but only a small one that fits into level-1 cache and a segmented search should be used). Now the optimum value has to be found from which on an event is considered rare. Very surprisingly, it turns out that the search is fastest if *all* events are considered frequent! This means we can forget about the priority queues. A better suited algorithm (and implementation) for a priority queue might give different results.

The resulting routine is remarkably fast, it uses just slightly more than 11 cycles per test. It was used to determine all candidates $x \le 10^{14}$. The search took about 8 days. The last entries in the list of candidates are

```
99205431802196^2+1 = [13.29.37.53.89.157.241.257.337.373.401^2.761]
99238108604548^2+1 = [5.29.37.61^2.101.349.397^2.433.557^2.661]
99311314035643^2+1 = [2.5^2.13.29.73.113.233.241.269.281.293.317.349.461]
99395767528881^2+1 = [2.13.29.37.53.149.173.181.193.313.353.373.401.449]
99501239756693^2+1 = [2.5^4.13.29.37^2.61.233.277.313^3.317.401]
99627378461772^2+1 = [5.13^2.37.41.73^2.137.277.281.521.557.617.761]
99759820688082^2+1 = [5^2.17.29.37.109.181.257.269.337.389.409.457.653]
99849755159917^2+1 = [2.5.89.101.181.233.257.293.389.457.521.557.677]
99950583525307^2+1 = [2.5^3.13.173.181.193.241^3.257.457^2.677]
99955223464153^2+1 = [2.5.13.61^2.101^2.109.373.421.433.509.709.757]
```

The search produced 43,936 candidates (including 0 and 1). Exactly that many logarithms were computed. This means that on average one logarithm was computed for one in $10^{14}/43,936 > 2 \cdot 10^9$ values tested.

One can further extend the list by testing (for each element x) whether the right hand side formula

```
[x] == [x+d] + [x+(x^2+1)/d]  where d divides x^2+1
```

leads to new candidate $x+d$ and $x+(x^2+1)/d$. Additionally one can try the arguments on the right hand side of relations like

```
[x] == 2[2*x] - [4*x^3+3*x]
[x] == [2*x-1] + [2*x+1] - [2*x^3+x]
[x] == 3[3*x] - [(9*x^3+7*x)/2] - [(27*x^3+9*x)/2]
```

Michael Roby Wetherfield has developed a more sophisticated approach for extending the list and sent me a big set of candidates beyond 10^{14}. His methods are described in [271] (see also [272], [253], [35], and [179]). We note that a single value, $x = 276,914,859,479,857,813,947$ where

```
x^2+1 = [2.5.13.17.29^3.41.53^2.73^2.101.157.181.229.241.313.397.401.509.577]
```

was discarded because it is bigger than $2^{64} = 18,446,744,073,709,551,616$.

We note the curious relation

$$[k\,a] \quad = \quad [(k+1)\,a] + [(k+1)\,k\,a] - [(k^4 + 2\,k^3 + k^2)\,a^3 + (k^2 + k + 1)\,a] \tag{30.5-5}$$

Set $f(a,k) := (k^4 + 2\,k^3 + k^2)\,a^3 + (k^2 + k + 1)\,a$, then

$$f(a,k)^2 + 1 \quad = \quad \Big((k\,a)^2 + 1\Big) \cdot \Big(((k+1)\,a)^2 + 1\Big) \cdot \Big(((k+1)\,k\,a)^2 + 1\Big) \tag{30.5-6}$$

30.5.4 Searching for all n-term relations

To find all n-term relations whose arguments are a subset of our just determined list of candidates, we have to test all subsets of $(n-1)$ (out of 64) odd primes, select the corresponding values x, and compute the nullspace as described. Let A_j be the j-th candidate. An array M of 64-bit auxiliary values is used. Its j-th entry M_j is a bit-mask corresponding to the odd primes in the factorization of $A_j^2 + 1$: bit i of M_j is set if the i-th odd prime divides $A_j^2 + 1$.

To find n-term relations, we must try all $\binom{64}{n-1}$ subsets of size $n-1$ out of the 64 odd primes in our scope. The bit-combination routine from section 1.26 on page 66 was used for this task. The selection of the entries that factor completely in the subset of $n-1$ primes under consideration can be done with a single bit-AND and a branch. The candidates with more than $n-1$ odd primes in their factorization should be discarded before the search.

While the search is very fast for small n, it does not finish in reasonable time for $n > 8$. A considerable speedup can be achieved by splitting our $N = 64$ odd primes into a group of the 20 smallest and $b = 64 - 20 = 44$ 'big' primes. Write ($q = n - 1$ and)

$$\binom{N}{q} \quad = \quad \binom{b}{0}\binom{N-b}{q} + \binom{b}{1}\binom{N-b}{q-1} + \binom{b}{2}\binom{N-b}{q-2} + \ldots \tag{30.5-7a}$$

$$\quad = \quad \sum_{j=0}^{q} \binom{b}{j}\binom{N-b}{q-j} \tag{30.5-7b}$$

This means, we first select the $j = 0,\ 1,\ 2,\ \ldots$-subsets of the big primes. We copy the corresponding candidates whose big prime factors are in the current subset into a new array B. The size of B will be significantly smaller than the size of A. From this array we select the arguments according to subsets of the small primes (leaving the subset of big primes fixed). This results in a much improved memory locality and accelerates the search by a factor of about 25.

Still, the limit for n so that an exhaustive search can be done has only been moved a little. But if we look at the prime sets that lead to the best relations, we observe that small primes are much 'preferred':

```
n     prime set of best relation
2     {13}
3     {5, 13}
4     {5, 13, 61}
5     {5, 13, 73, 101}
6     {5, 13, 61, 89, 197}
7     {5, 13, 17, 29, 97, 433}
8     {5, 13, 29, 37, 61, 97, 337}
9     {5, 13, 17, 29, 41, 53, 97, 269}
10    {5, 13, 17, 41, 53, 73, 97, 101, 157}
11    {5, 13, 17, 29, 37, 53, 61, 89, 97, 101}
12    {5, 13, 17, 29, 37, 53, 61, 89, 97, 101, 197}
13    {5, 13, 17, 29, 37, 53, 61, 89, 97, 101, 181, 281}
14    {5, 13, 17, 29, 37, 53, 61, 89, 97, 101, 181, 269, 457}
15    {5, 13, 17, 29, 37, 41, 53, 61, 89, 97, 101, 181, 337, 389}
```

The data suggests that the best possible relation is found long before the search space is exhausted. Therefore we will stop after the number of big primes in the subset is greater than, say, 4. In practice both parameters, the number b of primes considered big and the maximum number of primes taken from that set should be chosen depending on n.

Another important improvement is to discard small candidates before the search. One thereby avoids the huge amount of uninteresting relations with small first arguments x_1. Obviously, the amount of nullspace computation is also reduced significantly.

The results of the searches can be found in [18]. While the searches for the n-term relations with $n > 11$ did not even exhaust the table of candidates (which in turn is incomplete!) one can be reasonably sure that the best relations within our scope (of the first 64 odd primes $4\,i + 1$) have been found. Indeed I do not expect to see a better relation for any $n \leq 15$.

To improve on the results, one should use the first 128 odd primes $4\,i + 1$, sieve up to 10^{16} (distributed on 100 machines), and use a 3-phase subset selection instead of the described 2-phase selection. The selection (an nullspace computation) stage should also be done in a distributed fashion to reasonably exhaust the table of candidates. Such a computation would likely improve on some of the relations with more than 17 terms and produce up to 35-term relations that are in the vicinity of the best possible.

A method for the simultaneous computation of logarithms of small primes that uses a similar method to the one given here is described in section 31.4 on page 628.

Chapter 31

Logarithm and exponential function

We describe algorithms for the computation of the exponential (and hyperbolic cosine) function, and the logarithm (and inverse tangent). Constructions of superlinear iterations the compute the functions from their inverses are given.

We give argument reduction schemes and methods for the fast computation of the exponential and logarithm of power series.

31.1 Logarithm

31.1.1 AGM-based computation

The (natural) logarithm can be computed using the relation (see [56, p.221])

$$\left|\log(d) - R'(10^{-n}) + R'(10^{-n}\,d)\right| \;\leq\; \frac{n}{10^{2(n-1)}} \tag{31.1-1a}$$

$$\log(d) \;\approx\; R'(10^{-n}) - R'(10^{-n}\,d) \tag{31.1-1b}$$

which holds for $n \geq 3$ and $d \in]\frac{1}{2}, 1[$. Note that the first term on the right hand side is constant and can be saved for subsequent log-computations. One uses the relation

$$\log(M\,R^X) \;=\; \log(M) + X\,\log(R) \tag{31.1-2}$$

where M is the mantissa, R the radix, and X the exponent of the floating point representation. The value $\log(R)$ is computed only once. If M is not in the interval [1/2, 3/2] an argument reduction is made via

$$\log(M) \;=\; \log\left(M\,s^f\right) - f\,\log(s) \tag{31.1-3}$$

Where $0 < M < 1$ for the mantissa M, $s = \sqrt{2}$, and $f \in \mathbb{Z}$ so that $M\,s^f \in [1/2,\, 3/2]$. The quantity $\log(s) = \log(\sqrt{2})$ can be precomputed directly via the AGM. A C++ implementation is given in [hfloat: src/tz/log.cc]:

There is a nice way to compute the value of $\log(R)$ if the value of π has been precomputed. One defines

$$\Theta_2(q) = \sum_{n=-\infty}^{\infty} q^{(n+1/2)^2} = 2\left(q^{1/4} + q^{9/4} + q^{25/4} + q^{81/4} + q^{121/4} + \ldots\right) \tag{31.1-4a}$$

$$= 2\,q^{1/4}\left(1 + q^2 + q^6 + q^{12} + q^{20} + \ldots\right) \tag{31.1-4b}$$

$$\Theta_3(q) = \sum_{n=-\infty}^{\infty} q^{n^2} = 1 + 2\left(q + q^4 + q^9 + q^{16} + q^{25} + \ldots\right) \tag{31.1-4c}$$

$$\Theta_4(q) = \sum_{n=-\infty}^{\infty} (-1)^n\, q^{n^2} = 1 + 2\left(-q + q^4 - q^9 + q^{16} - q^{25} \pm \ldots\right) \tag{31.1-4d}$$

Then (with K the elliptic function, see section 30.2.1 on page 595)

$$\log\frac{1}{q} = -\log q = \pi\,\frac{K'}{K} \tag{31.1-5}$$

thereby $q = \exp(-\pi\,K'/K)$, and

$$\frac{\pi}{\log(1/q)} = -\frac{\pi}{\log(q)} = \mathrm{AGM}\left(\Theta_3^2(q),\,\Theta_2^2(q)\right) = \frac{\mathrm{AGM}(1,\,k)}{\mathrm{AGM}(1,\,k')} \tag{31.1-6}$$

Computing $\Theta_3(q)$ is easy if $q = 1/R$:

$$\Theta_3(q) = 1 + 2\sum_{n=1}^{\infty} q^{n^2} = 2\,(1 + \sum_{n=1}^{\infty} q^{n^2}) - 1 \tag{31.1-7}$$

However, the computation of $\Theta_2(q)$ suggests to choose $q = 1/R^4 =: b^4$:

$$\Theta_2(q) = 0 + 2\sum_{n=0}^{\infty} q^{(n+1/2)^2} = 2\sum_{n=0}^{\infty} b^{4n^2+4n+1} \quad \text{where} \quad q = b^4 \tag{31.1-8a}$$

$$= 2\,b\sum_{n=0}^{\infty} q^{n^2+n} = 2\,b\,(1 + \sum_{n=1}^{\infty} q^{n^2+n}) \tag{31.1-8b}$$

Functions that compute $\Theta_2(b^4)$, $\Theta_3(b^4)$ and $\pi/\log(b)$ where b is the inverse fourth power of the used radix R are given in [hfloat: src/tz/pilogq.cc].

We list a few relations involving the theta functions (see, for example [273]):

$$\Theta_3^4(q) = \Theta_2^4(q) + \Theta_4^4(q) \tag{31.1-9}$$

$$\Theta_3(q) = \Theta_3(q^4) + \Theta_2(q^4) \qquad \Theta_4(q) = \Theta_3(q^4) - \Theta_2(q^4) \tag{31.1-10}$$

$$\Theta_2^2(q) = \frac{2\,k\,K}{\pi} \qquad \Theta_3^2(q) = \frac{2\,K}{\pi} \qquad \Theta_4^2(q) = \frac{2\,k'\,K}{\pi} \tag{31.1-11}$$

$$k = \frac{\Theta_2^2(q)}{\Theta_3^2(q)} \qquad k' = \frac{\Theta_4^2(q)}{\Theta_3^2(q)} \tag{31.1-12}$$

$$1 = \mathrm{AGM}\left(\Theta_3^2(q), \Theta_4^2(q)\right) \tag{31.1-13}$$

$$q\prod_{n=1}^{\infty}(1-q^n)^{24} = \frac{256}{\pi^{12}}\,k^2\,k'^8\,K^{12} \tag{31.1-14}$$

31.1.2 Computation by inverting the exponential function

31.1.2.1 Iterations from the power series

With an efficient algorithm for the exponential function one can compute the logarithm using

$$y := 1 - d\,e^{-x} \tag{31.1-15}$$

$$\log(d) = x + \log(1-y) \tag{31.1-16}$$

$$= x + \log\left(1-(1-d\,e^{-x})\right) = x + \log\left(e^{-x}d\right) = x + (-x + \log(d)) \tag{31.1-17}$$

Thereby,

$$\log(d) = x + \log(1-y) = x - \left(y + \frac{y^2}{2} + \frac{y^3}{3} + \frac{y^4}{4} + \dots\right) \tag{31.1-18}$$

Truncation of the series before the n-th power of y gives an iteration of order n:

$$x_{k+1} = \Phi_n(x_k) := x - \left(y + \frac{y^2}{2} + \frac{y^3}{3} + \cdots + \frac{y^{n-1}}{n-1}\right) \tag{31.1-19}$$

31.1.2.2 Iterations from Padé approximants

Padé approximants $P_{[i,j]}(z)$ of $\log(1-z)$ at $z=0$ produce iterations of order $i+j+1$. Compared to the power series based iteration one needs one additional long division but saves half of the exponentiations. This can be a substantial saving for high order iterations.

The approximants can be computed via the continued fraction expansion of $\log(1+z)$:

$$\log(1+z) = 0 + \cfrac{c_1\,z}{1 + \cfrac{c_2\,z}{1 + \cfrac{c_3\,z}{1 + \cfrac{c_4\,z}{1+\cdots}}}} \tag{31.1-20}$$

where $c_1 = 1$ and

$$c_k = \frac{k}{4(k-1)} \quad \text{if } k \text{ even}, \qquad c_k = \frac{k-1}{4k} \quad \text{else} \tag{31.1-21}$$

Using recurrence relations 35.3-7a and 35.3-7b on page 704 with $a_0 = 0$, $a_k = 1$ and $b_k = c_k \cdot z$ we obtain (fractions in lowest terms):

$$1 \mapsto P_{[1,0]} = \frac{z}{1} \tag{31.1-22a}$$

$$2 \mapsto P_{[1,1]} = \frac{2\,z}{2+\,z} \tag{31.1-22b}$$

$$3 \mapsto P_{[2,1]} = \frac{6\,z+\,z^2}{6+4\,z} \tag{31.1-22c}$$

$$4 \mapsto P_{[2,2]} = \frac{6\,z+3\,z^2}{6+6\,z+\,z^2} \tag{31.1-22d}$$

$$5 \mapsto P_{[3,2]} = \frac{30\,z+21\,z^2+\,z^3}{30+36\,z+9\,z^2} \tag{31.1-22e}$$

$$6 \mapsto P_{[3,3]} = \frac{60\,z+60\,z^2+11\,z^3}{60+90\,z+36\,z^2+3\,z^3} \tag{31.1-22f}$$

$$7 \mapsto P_{[4,3]} = \frac{420\,z+510\,z^2+140\,z^3+3\,z^4}{420+720\,z+360\,z^2+48\,z^3} \tag{31.1-22g}$$

$$8 \mapsto P_{[4,4]} = \frac{420\,z+630\,z^2+260\,z^3+25\,z^4}{420+840\,z+540\,z^2+120\,z^3+6\,z^4} \tag{31.1-22h}$$

The expressions are Padé approximants correct up to order k. For even k these are the diagonal approximants $[k/2, k/2]$ which satisfy the functional equation $\log(1/z) = -\log(z)$: $P(1/z-1) = -P(z-1)$. Further information like the error term of the diagonal approximants is given in [189].

The diagonal Padé approximants can be computed by setting $P_0 = 0$, $Q_0 = 1$, $P_2 = z$, $Q_2 = 1 + z/2$, and computing, for $k = 4, 6, \ldots, 2n$,

$$P_k = A_k P_{k-2} + B_k P_{k-4} \tag{31.1-23a}$$
$$Q_k = A_k Q_{k-2} + B_k Q_{k-4} \tag{31.1-23b}$$

(these are relations 35.3-14a and 35.3-14b on page 708). The A_k and B_k are defined as

$$A_k = 1 + z/2 \tag{31.1-23c}$$
$$B_k = \frac{z^2}{16} \frac{(k-2)^2}{1-(k-2)^2} \tag{31.1-23d}$$

Then P_{2n}/Q_{2n} is the Padé approximant $[n, n]$ of $\log(1+z)$ which is correct up to order $2\,n$. The following pari/gp function implements the algorithm:

```
1   log_pade(n, z='z)=
2   { /* Return Pade approximant [n,n] of log(1+z) */
3       local(P0,Q0,P2,Q2,tp,tq, t);
4       if ( n<1, return(0) );
5       P0=0;  Q0=1;
6       P2=z;  Q2=1+z/2;
7       forstep (k=4, 2*n, 2,
8           Ak = 1+z/2;  \\ == +z*C(k-1)+z*C(k)+1;
9           t = (k-2)^2;
10          Bk = z^2/16*t/(1-t);  \\ == -z^2*C(k-1)*C(k-2);
11          tp = Ak*P2 + Bk*P0;
12          tq = Ak*Q2 + Bk*Q0;
13          P0=P2; P2=tp;
14          Q0=Q2; Q2=tq;
15      );
16      return( P2/Q2 );
17  }
```

31.1.2.3 Padé approximants for arctan *

A continued fraction for arctan is (given in [188, p.569])

$$\arctan(z) = z \cfrac{1}{1 + \cfrac{z^2/(1\cdot 3)}{1 + \cfrac{2^2 z^2/(3\cdot 5)}{1 + \cfrac{3^2 z^2/(5\cdot 7)}{1 + \cdots}}}} \tag{31.1-24}$$

Padé approximants P_k/Q_k for the inverse tangent can be obtained by setting $P_0 = 0$, $P_1 = z$, $Q_0 = 1$, $Q_1 = 1$, and the recurrences

$$P_{k+1} = P_k + P_{k-1} \frac{k^2 z^2}{4k^2 - 1} \tag{31.1-25a}$$
$$Q_{k+1} = Q_k + P_{k-1} \frac{k^2 z^2}{4k^2 - 1} \tag{31.1-25b}$$

The first few approximants are

$$2 \mapsto P_{[1,2]} = \frac{3\,x}{3+x^2} \tag{31.1-26a}$$

$$3 \mapsto P_{[3,2]} = \frac{15\,x+4\,x^3}{15+9\,x^2} \tag{31.1-26b}$$

$$4 \mapsto P_{[3,4]} = \frac{105\,x+55\,x^3}{105+90\,x^2+9\,x^4} \tag{31.1-26c}$$

$$5 \mapsto P_{[5,4]} = \frac{945\,x+735\,x^3+64\,x^5}{945+1050\,x^2+225\,x^4} \tag{31.1-26d}$$

$$6 \mapsto P_{[5,6]} = \frac{1155\,x+1190\,x^3+231\,x^5}{1155+1575\,x^2+525\,x^4+25\,x^6} \tag{31.1-26e}$$

$$k \mapsto P_{[i,j]} = \arctan(z)+O(z^{2k+1}) \tag{31.1-26f}$$

31.1.3 Argument reduction for the logarithm

When the logarithm shall be computed for high but not extreme precision (up to several hundred decimal digits or so), the following scheme can beat the AGM algorithm. Use the functional equation for the logarithm

$$\log(z^a) = a\,\log(z) \tag{31.1-27}$$

to reduce the argument by setting $a = 1/N$. Now with N big enough $z^{1/N}$ will be close to one: $r := z^{1/N} = 1+e$ where e is small. Then a few terms of the Taylor series of $\log(1+e) = e - e^2/2 + e^3/3 \pm \ldots$ suffice to compute the logarithm. Compute the logarithm of z as follows

1. set $r = z^{1/N}$ and $e = r-1$
2. compute $l := \log(1+e)$ to the desired precision using the Taylor series
3. return $L := N\,l$

One can also use a Padé approximant in step 2. With argument $z = 2.0$, $N = 2^{32}$, and four terms of the Taylor series we obtain:

```
? z=2.0;   \\ argument for log()
? n=32; N=2^n;
? r=z^(1/N)  \\ compute by 32 sqrt extractions
  1.00000000016138590420965976120397663110198503274461201650532 65785 E-10
? e=r-1; \\ small
  1.61385904209659761203976631101985032744612016505326578 5 E-10
? l=e-1/2*e^2+1/3*e^3-1/4*e^4  \\ approx log(1+e)
  1.61385904196637056166593013670802248659405413369314055 0 E-10
? L=N*l     \\ final result
  0.693147180559945309417232121458176568075406093226565036 5
? log(z)  \\ check with builtin log
  0.693147180559945309417232121458176568075500134360255254 1
```

One may also use the following reduction for $L(z) := \log(1+z)$, which avoids loss of precision for small values of z:

$$L(z) = 2\,L\left(\frac{z}{1+\sqrt{1+z}}\right) \tag{31.1-28}$$

31.1.4 Argument reduction for arctan

We use the equation

$$\arctan(z) = 2\,\arctan\left(\frac{z}{1+\sqrt{1+z^2}}\right) \tag{31.1-29}$$

Compute the inverse tangent of z as follows

1. set $r := z$
2. repeat n times: $r = r/(1+\sqrt{1+r^2})$ (for n big enough)
3. compute $a := \arctan(r)$ to the desired precision using the Taylor series
4. return $A := 2^n\, a$

We compute arctan(1.0) using $n = 16$ and four terms of the Taylor series:

```
? z=1.0;
? n=16;
? r=z;for(k=1,n,r=r/(1+sqrt(1+r^2)));r
  0.000011984224905930303478511634794650661319585128744025261 89
? a=r-1/3*r^3+1/5*r^5-1/7*r^7
  0.000011984224905356572107172559292905818497455676569676602 63
? A=2^n*a
  0.785398163397448309615660845819875721049255219670325829 9
? atan(z)  \\ check with builtin atan
  0.785398163397448309615660845819875721049292349843776455 2
```

All divisions in the reduction phase can be saved by using

$$\arctan(1/z) \quad = \quad 2\,\arctan\left(\frac{1}{1+\sqrt{1+z^2}}\right) \tag{31.1-30}$$

The inverse sine and cosine can be computed as

$$\arcsin(z) \quad = \quad \arctan\left(\frac{z}{\sqrt{1-z^2}}\right) \tag{31.1-31a}$$

$$\arccos(z) \quad = \quad \arctan\left(\frac{\sqrt{1-z^2}}{z}\right) \tag{31.1-31b}$$

31.1.5 A curious series for the logarithm *

We finally note two relations resembling the well-known series

$$\frac{1}{2}\,\log\left(\frac{1+x}{1-x}\right) \quad = \quad x + \frac{1}{3}\,x^3 + \frac{1}{5}\,x^5 + \ldots + \frac{1}{2k+1}\,x^{2k+1} + \ldots \tag{31.1-32}$$

The first is

$$\frac{1}{6}\,\log\left(\frac{1+3x+3x^2}{1-3x+3x^2}\right) \quad = \quad x - \frac{3^2}{5}\,x^5 - \frac{3^3}{7}\,x^7 + \frac{3^5}{11}\,x^{11} + \frac{3^6}{13}\,x^{13} \pm \ldots = \tag{31.1-33a}$$

$$= \quad \sum_{k=0}^{\infty}\left(+\frac{3^{6k+0}}{12k+1}\,x^{12k+1} - \frac{3^{6k+2}}{12k+5}\,x^{12k+5} - \frac{3^{6k+3}}{12k+7}\,x^{12k+7} + \frac{3^{6k+5}}{12k+11}\,x^{12k+11}\right) \tag{31.1-33b}$$

The second, given in [26], is

$$\frac{1}{3}\,\log\left(\frac{1+x+x^2}{1-2\,x+x^2}\right) \quad = \quad x + \frac{x^2}{2} + \frac{x^4}{4} + \frac{x^5}{5} + \ldots \;=\; \sum_{k=0}^{\infty}\left(+\frac{x^{3k+1}}{3k+1} + \frac{x^{3k+2}}{3k+2}\right) \tag{31.1-34}$$

Relation 31.1-33a can be brought in a similar form, namely

$$\frac{1}{2\sqrt{3}}\,\log\left(\frac{1+\sqrt{3}\,x+x^2}{1-\sqrt{3}\,x+x^2}\right) \quad = \quad x - \frac{x^5}{5} - \frac{x^7}{7} + \frac{x^{11}}{11} + \frac{x^{13}}{13} \pm \ldots = \tag{31.1-35a}$$

$$= \quad \sum_{k=0}^{\infty}\left(+\frac{x^{12k+1}}{12k+1} - \frac{x^{12k+5}}{12k+5} - \frac{x^{12k+7}}{12k+7} + \frac{x^{12k+11}}{12k+11}\right) \tag{31.1-35b}$$

31.2 Exponential function

31.2.1 AGM-based computation of the exponential function

We use the following relation (see [163, pp.35-36]):

$$q \;=\; \exp\left(-\pi\,\frac{K'}{K}\right) \tag{31.2-1}$$

Now write

$$\frac{K'}{K} \;=\; \frac{\mathrm{AGM}(1,k')}{\mathrm{AGM}(1,k)} \;=\; \frac{\mathrm{AGM}(1,b_0)}{\mathrm{AGM}(1,b'_0)} \tag{31.2-2}$$

where $k' = b_0$ and $k = b'_0 = \sqrt{1-b_0^2}$ and use ([163, p.38], note the missing '4' there)

$$\frac{\pi}{2}\,\frac{\mathrm{AGM}(1,b_0)}{\mathrm{AGM}(1,b'_0)} \;=\; \lim_{n\to\infty}\frac{1}{2^n}\log\frac{4\,a_n}{c_n} \tag{31.2-3}$$

thereby

$$q \;=\; \exp\left(-2\,\lim_{n\to\infty}\frac{1}{2^n}\,\log\frac{4\,a_n}{c_n}\right) \;=\; \lim_{n\to\infty}\exp\left(-\frac{1}{2^{n-1}}\,\log\frac{4\,a_n}{c_n}\right) \tag{31.2-4a}$$

$$\;=\; \lim_{n\to\infty}\left(\exp\log\frac{4\,a_n}{c_n}\right)^{-1/2^{n-1}} \;=\; \lim_{n\to\infty}\left(\frac{4\,a_n}{c_n}\right)^{-1/(2^{n-1})} \tag{31.2-4b}$$

This gives

$$q \;=\; \lim_{n\to\infty}\left(\frac{c_n}{4\,a_n}\right)^{1/(2^{n-1})} \tag{31.2-5}$$

One obtains an algorithm for $\exp(-x)$ by first solving for k, k' so that $x = \pi\,K'/K$ (precomputed π) and applying the last relation that implies the computation of a 2^{n-1}-th root. Note that the quantity c should be computed via $c_{n+1} = \frac{c_n^2}{4\,a_{n+1}}$ throughout the AGM computation in order to preserve its accuracy.

For $k = 1/\sqrt{2} =: s$ one has $k = k'$ and so $q = \exp(-\pi)$. Thus the calculation of $\exp(-\pi) = 0.0432139182637\ldots$ can directly be done via a single AGM computation as $(c_n/(4a_n))^N$ where $N = 1/2^{(n-1)}$. The quantity $i^i = \exp(-\pi/2) = 0.2078795763507\ldots$ can be obtained using $N = 1/2^n$.

31.2.2 Computation by inverting the logarithm

31.2.2.1 Iterations from the power series

The exponential function can be computed using the iteration that is obtained as follows:

$$\exp(d) \;=\; x\,\exp\left(d-\log(x)\right) \tag{31.2-6}$$

$$\;=\; x\,\exp(y) \quad \text{where} \quad y := d-\log(x) \tag{31.2-7}$$

$$\;=\; x\left(1+y+\frac{y^2}{2}+\frac{y^3}{3!}+\frac{y^4}{4!}+\ldots\right) \tag{31.2-8}$$

The corresponding n-th oder iteration is

$$x_{k+1} \;=\; \Phi_n(x_k) \;:=\; x_k\left(1+y+\frac{y^2}{2}+\frac{y^3}{3!}+\ldots+\frac{y^{n-1}}{(n-1)!}\right) \tag{31.2-9}$$

As the computation of logarithms is expensive one should use an iteration of high order. The C++ implementation given in [hfloat: src/tz/itexp.cc] uses the iteration of order 20.

31.2.2.2 Iterations from Padé approximants

Padé approximants $P_{[i,j]}(z)$ of $\exp(z)$ around $z=0$ give iterations of order $i+j+1$. For $i=j$ we obtain

$$P_{[1,1]} = \frac{2+z}{2-z} \tag{31.2-10a}$$

$$P_{[2,2]} = \frac{12+6z+z^2}{12-6z+z^2} \tag{31.2-10b}$$

$$P_{[3,3]} = \frac{120+60z+12z^2+z^3}{120-60z+12z^2-z^3} \tag{31.2-10c}$$

$$P_{[4,4]} = \frac{1680+840z+180z^2+20z^3+z^4}{1680-840z+180z^2-20z^3+z^4} \tag{31.2-10d}$$

Note that the functional equation $\exp(-z)=\frac{1}{\exp(z)}$ holds for the diagonal approximants. In general, we have $P_{[i,j]}(-z)=1/P_{[j,i]}(z)$. This can be seen from the following closed form for the Padé approximants:

$$P_{[i,j]}(z) = \left\{\sum_{k=0}^{i} \frac{\binom{i}{k}}{\binom{i+j}{k}} \frac{z^k}{k!}\right\} / \left\{\sum_{k=0}^{j} \frac{\binom{j}{k}}{\binom{i+j}{k}} \frac{(-z)^k}{k!}\right\} \tag{31.2-11}$$

The numerator for $i=j$, multiplied by $(2i)!/i!$ in order to avoid rational coefficients, equals

$$= \frac{(2i)!}{i!} \cdot \sum_{k=0}^{i} \frac{\binom{i}{k}}{\binom{2i}{k}} \frac{z^k}{k!} \tag{31.2-12}$$

The coefficients of the numerator and denominator in the diagonal approximant

$$P_{[i,i]} = \frac{\sum_{k=0}^{i} c_k\, z^k}{\sum_{k=0}^{i} c_k\, (-z)^k} \tag{31.2-13}$$

can be computed using $c_i=1$ (the coefficient of the highest power of z) and the recurrence

$$c_k = c_{k+1}\, \frac{(k+1)\,(2i-k)}{(i-k)} \tag{31.2-14}$$

It is usually preferable to generate the coefficients in the other direction. To do so compute the constant c_0

$$c_0 = \prod_{w=1}^{i} 4\,w-2 \;=\; 2,\ 12,\ 120,\ 1680,\ 30240,\ \ldots \tag{31.2-15}$$

and use the recurrence

$$c_k = c_{k-1}\, \frac{i-k}{(2i-k)\,(k+1)} \tag{31.2-16}$$

We generate the coefficients for $1 \le i \le 8$:

```
? c0(i)=prod(w=1,i,4*w-2)
? qq(i,k)=(i-k)/((2*i-k)*(k+1))
? for (i=1, 8, c=c0(i); print1("["i,","i,"]:  "); \
?   for (k=0, i, print1("  ", c); c*=qq(i,k)); print();)
  [1,1]:    2  1
  [2,2]:    12  6  1
  [3,3]:    120  60  12  1
  [4,4]:    1680  840  180  20  1
  [5,5]:    30240  15120  3360  420  30  1
  [6,6]:    665280  332640  75600  10080  840  42  1
  [7,7]:    17297280  8648640  1995840  277200  25200  1512  56  1
  [8,8]:    518918400  259459200  60540480  8648640  831600  55440  2520  72  1
```

Finally, the Padé approximant $P_{[i,j]}$ can be expressed as ratio of hypergeometric functions:

$$P_{[i,j]} \quad = \quad F\left(\begin{matrix}-i\\-i-j\end{matrix}\Big|\, z\right) / F\left(\begin{matrix}-j\\-i-j\end{matrix}\Big|\, -z\right) \tag{31.2-17}$$

This is relation 35.2-53 on page 693 with $a = -i$ and $b = -j$ where i and j are positive integers.

31.2.3 Argument reduction for the exponential function

As for the logarithm an argument reduction technique can be useful with moderate precisions. We do not use the functional equation for the exponential functional ($\exp(2z) = \exp(z)^2$) because the loss of precision when adding up the Taylor series (one plus a tiny quantity). Instead we use the functional equation for $E(z) := \exp(z) - 1$:

$$E(2z) \quad = \quad 2\,E(z) + E^2(z) \tag{31.2-18}$$

Compute the exponential function of z as follows

1. set $r = z/2^n$ (for n big enough)
2. compute $E := \exp(r) - 1$ to the desired precision using the Taylor series
3. repeat n times: $E = 2\,E + E^2$
4. return $E + 1$

We compute exp(1.0) using $n = 16$ and eight terms of the Taylor series:

```
? z=1.0;
? n=16;
? r=z/2^n
  0.000015258789062500000000000000000000000000000000000000000000000
? E=r*(1+r/2*(1+r/3*(1+r/4*(1+r/5*(1+r/6*(1+r/7*(1+r/8)))))))
  0.000015258905478413948140042622480661730187012348455116225 83
? for(k=1,n,E=2*E+E^2);E=E+1
  2.718281828459045235360287471352662497757247071686614582
? exp(1.0)  \\ check with builtin exp()
  2.718281828459045235360287471352662497757247093699959575
```

One can also compute the exponential function via the hyperbolic cosine using

$$\exp(z) \quad = \quad \cosh(z) + \sinh(z) \;=\; \cosh(z) + \sqrt{\cosh^2(z) - 1} \tag{31.2-19a}$$

Alternatively, compute hyperbolic sine and use

$$\exp(z) \quad = \quad \sinh(z) + \sqrt{1 + \sinh(z)^2} \tag{31.2-19b}$$

The advantage is that half of the coefficients of the Taylor series are zero. Again we do not use the function equation for the cosine ($\cosh(2z) = 2\,\cosh^2(z) - 1$) but that for $C(z) := \cosh(z) - 1$:

$$C(2z) \quad = \quad 2\,(C(z) + 1)^2 - 2 \tag{31.2-20}$$

Compute the hyperbolic cosine as follows

1. set $r = z/2^n$ (for n big enough)
2. compute $C := \cosh(r) - 1$ to the desired precision using the Taylor series
3. repeat n times: $C = 2\,[C - 1]^2 - 2$
4. return $C + 1$

We compute cosh(1.5) using $n = 16$ and four terms of the Taylor series:

```
? z=1.50;
? n=16;
? r=z/2^n
? C=1/2*r^2+1/24*r^4+1/720*r^6+1/40320*r^8
  0.000000000261934474122038277307663984994058290343310054242645592 3
? for(k=1,n,C=2*(C+1)^2-2);C
  1.352409615243247325767667965441644170173960682574839216
? C+1
  2.352409615243247325767667965441644170173960682574839216
? cosh(z) \\ check with builtin cosh
  2.352409615243247325767667965441644170173960748865373192
```

When the Taylor series for $\cos(z)-1$ is used then the cosine can be obtained by the identical algorithm:

```
? z=1.50;
? n=16;
? r=z/2^n
  0.000022888183593750000000000000000000000000000000000000000000
? C=-1/2*r^2+1/24*r^4-1/720*r^6+1/40320*r^8
  -0.000000000261934474099168387731797875839209846883179633343528818 1
? for(k=1,n,C=2*(C+1)^2-2);C
  -0.92926279833229708991181014856573129091490894138179316 23
? C+1
  0.070737201667702910088189851434268709085091058618206837 69
? cos(z)  \\ check with builtin cos
  0.070737201667702910088189851434268709085091027563346869 42
```

Compute the sine as $\sin(z)=\sqrt{1-\cos^2(z)}$, and the tangent as $\tan(z)=\sin(z)/\cos(z)$.

31.3 Logarithm and exponential function of power series

The computation of the logarithm, the exponential function, and the inverse trigonometric functions turns out to be surprisingly simple with power series.

31.3.1 Logarithm

Let $f(x)$ be a power series in x and $g(x)=\log(f(x))$. Then we have $\frac{d\,g(x)}{dx}=\frac{f'(x)}{f(x)}$. Thereby, symbolically,

$$g(x) \;=\; \log(f(x)) \;=\; \int \frac{f'(x)}{f(x)}\,dx \tag{31.3-1}$$

A few lines of pari/gp to demonstrate this:

```
? sp=8;default(seriesprecision,sp+1);
? f=taylor((1)/(1-x-x^2),x) /* shifted Fibonacci (with constant term) */
 1 + x + 2*x^2 + 3*x^3 + 5*x^4 + 8*x^5 + 13*x^6 + 21*x^7 + 34*x^8 + O(x^9)

? d=deriv(f,x)
 1 + 4*x + 9*x^2 + 20*x^3 + 40*x^4 + 78*x^5 + 147*x^6 + 272*x^7 + O(x^8)

? q=d/f  /* the only nontrivial computation */
 1 + 3*x + 4*x^2 + 7*x^3 + 11*x^4 + 18*x^5 + 29*x^6 + 47*x^7 + O(x^8)

? lf=intformal(q)
 x + 3/2*x^2 + 4/3*x^3 + 7/4*x^4 + 11/5*x^5 + 3*x^6 + 29/7*x^7 + 47/8*x^8 + O(x^9)

? f-exp(lf)  /* check with builtin exp() */
 O(x^9)
```

31.3.2 Inverse trigonometric functions

Now let $a(x)=\arctan(f(x))$. Then, symbolically,

$$a(x) \;=\; \int \frac{f'(x)}{1+f(x)^2}\,dx \tag{31.3-2}$$

Verification for the trivial case $f(x)=x$:

```
? sp=13;default(seriesprecision,sp+1);
? f=taylor(x,x)
 x + O(x^14)
? d=deriv(f,x)
 1 + O(x^13)
? q=d/(1+f^2)
 1 - x^2 + x^4 - x^6 + x^8 - x^10 + x^12 + O(x^13)
? af=intformal(q,x)
 x - 1/3*x^3 + 1/5*x^5 - 1/7*x^7 + 1/9*x^9 - 1/11*x^11 + 1/13*x^13 + O(x^14)
? f-tan(af)  /* check with builtin tan() */
 O(x^14)
```

For $s(x) = \arcsin(f(x))$ use

$$s(x) \quad = \quad \int \frac{f'(x)}{\sqrt{1-f(x)^2}}\, dx \tag{31.3-3}$$

31.3.3 Exponential function

With $e(x) = \exp(f(x))$ we can use a scheme similar to those shown in section 28.7 on page 577. We express a function $g(y)$ as

$$g(y) \quad = \quad \prod_{k=1}^{\infty} [1 + T(Y_k)] \tag{31.3-4}$$

where $Y_1 = y$, $Y_{k+1} = N(Y_k)$ and $1 + T(y)$ is the truncated Taylor series of g. A second order product is obtained by taking $1 + T(y) = 1 + y$ (the series of $\exp(y)$ truncated before the second term) and

$$N(y) \quad = \quad f^{-1}\left(\frac{f(y)}{1+T(y)}\right) \tag{31.3-5}$$

For $g(y) = \exp(y)$ one obtains $N(y) = y - \log(1+y)$ and

$$\exp(y) \quad = \quad [1+y] \cdot [1+y-\log(1+y)] \cdot [1+y-\log(1+y)-\log(1+y-\log(1+y))] \cdots \tag{31.3-6a}$$

$$= \quad \prod_{k=1}^{\infty} [1+Y_k] \tag{31.3-6b}$$

where $Y_1 = y = f(x)$ and $Y_{k+1} = Y_k - \log(1+Y_k)$. The product $\prod_{k=1}^{N}$ is correct up to order y^{2^N-1}. The computation involves $N-2$ logarithms and $N-1$ multiplications. Implementation in pari/gp:

```
1   texp(y, N=5)=
2   {
3       local(Y, e, t);
4       Y = y;  e = 1 + Y;
5       for (k=2, N,
6           t = deriv(1+Y,x)/(1+Y);
7           t = intformal(t);  \\ here: t = log(1+Y);
8           Y -= t;
9           e *= (1+Y);
10      );
11      return( e );
12  }
```

Check:

```
? f=taylor((x)/(1-x-x^2),x)
  x + x^2 + 2*x^3 + 3*x^4 + 5*x^5 + 8*x^6 + 13*x^7 + 21*x^8 + ...
? e=exp(f)  /* builtin exp() */
  1 + x + 3/2*x^2 + 19/6*x^3 + 145/24*x^4 + 467/40*x^5 + 16051/720*x^6 + ...
? t=texp(f,4);
? t-e
  -1/32768*x^16 - 35/98304*x^17 - ...
```

31.4 Simultaneous computation of logarithms of small primes

We describe a method to compute the logarithms of a given set of (small) primes simultaneously. We use a method similarly to the one for finding arctan-relations for π given in section 30.5 on page 610. We define

$$L(z) \quad := \quad 2\,\mathrm{arccoth}(z) \;=\; 2\sum_{k=0}^{\infty}\frac{1}{(2k+1)\,z^{2k+1}} \tag{31.4-1}$$

and note that (relation 35.2-85d on page 697)

$$\log(z) \quad = \quad 2\,\mathrm{arccoth}\,\frac{z+1}{z-1} \tag{31.4-2}$$

We will determine a set of relations that express the logarithm of a prime as linear combination of terms $L(X_i)$ where the X_i are large integers so that the series for L converges fast.

```
S = { 51744295, 170918749, 265326335, 287080366, 362074049, 587270881,
      617831551, 740512499, 831409151, 1752438401, 2151548801, 2470954914, 3222617399 }

2:  [ -1595639, -17569128, -8662593, -31112926, -13108464, -11209640,
      -12907342, +9745611, -1705229, -12058985, +4580610, +4775383, -12972664 ]

3:  [ -2529028, -27846409, -13729885, -49312821, -20776424, -17766859,
      -20457653, +15446428, -2702724, -19113039, +7260095, +7568803, -20561186 ]

5:  [ -3704959, -40794252, -20113918, -72241977, -30436911, -26027978,
      -29969920, +22628608, -3959419, -28000096, +10635847, +11088096, -30121593 ]

7:  [ -4479525, -49322778, -24318973, -87345026, -36800111, -31469438,
      -36235490, +27359389, -4787183, -33853851, +12859398, +13406195, -36418872 ]

11: [ -5520004, -60779197, -29967648, -107633040, -45347835, -38778983,
      -44652067, +33714275, -5899123, -41717234, +15846307, +16520111, -44878044 ]

13: [ -5904566, -65013499, -32055403, -115131507, -48507081, -41480597,
      -47762841, +36063046, -6310097, -44623547, +16950271, +17671017, -48004561 ]

17: [ -6522115, -71813158, -35408027, -127172929, -53580360, -45818987,
      -52758281, +39834823, -6970060, -49290653, +18723073, +19519201, -53025282 ]

19: [ -6778159, -74632382, -36798067, -132165454, -55683805, -47617738,
      -54829453, +41398649, -7243689, -51225694, +19458099, +20285481, -55106936 ]

23: [ -7217972, -79475039, -39185776, -140741248, -59296949, -50707501,
      -58387161, +44084875, -7713709, -54549566, +20720673, +21601741, -58682649 ]

29: [ -7751584, -85350490, -42082712, -151146003, -63680669, -54456218,
      -62703622, +47343993, -8283970, -58582320, +22252516, +23198720, -63020955 ]

31: [ -7905109, -87040909, -42916186, -154139543, -64941904, -55534757,
      -63945506, +48281670, -8448039, -59742579, +22693241, +23658185, -64269124 ]

37: [ -8312407, -91525553, -45127374, -162081337, -68287932, -58396097,
      -67240196, +50769306, -8883311, -62820720, +23862474, +24877135, -67580488 ]

41: [ -8548719, -94127517, -46410292, -166689119, -70229278, -60056229,
      -69151756, +52212618, -9135853, -64606639, +24540856, +25584363, -69501722 ]
```

Figure 31.4-A: Relations for the fast computation of the logarithms of the primes up to 41.

Compute $\log(p_i)$ for the primes p_i in a predefined set P of n primes as follows:

1. Find a set S of numbers $X \in \mathbb{Z}$ so that X^2-1 factor completely into the primes in P.
2. Select a subset S of n (large) numbers X_k so that all $L(X_k)$ are linearly independent.
3. Try to find, for each prime p_i, a relation $\log(p_i) = \sum_{j=1}^{n} m_j\,L(X_j)$. If this fails return to step 2.

For example, with the first 13 primes ($P = \{2, 3, 5, 7, 11, \ldots, 41\}$) one can find

$$\begin{aligned} S \quad = \quad & \{X_1, X_2, \ldots, X_{13}\} \;= \\ & \{51744295, 170918749, 265326335, 287080366, 362074049, 587270881, \\ & 617831551, 740512499, 831409151, 1752438401, 2151548801, 2470954914, 3222617399\} \end{aligned} \tag{31.4-3}$$

We use the short form `p: [m1, m2, m3, ..., m13]` to denote a relation

$$\log(p) \quad = \quad \sum_{j=1}^{13} m_j \, L(X_j) \tag{31.4-4}$$

Now we have the relations given in figure 31.4-A, the first one is

$$\log(2) \quad = \quad -1595639 \, L(51,744,295) - 17569128 \, L(170,918,749) \pm \ldots - 12972664 \, L(3,222,617,399)$$

Note that the series with slowest convergence already gives more than 15 digits per term ($\log_{10}(51,744,295^2) \approx 15.4$), and the last series gives 19 digits per term.

```
                  2,  3,  5,  7, 11, 13, 17, 19, 23, 29, 31, 37, 41

   51744295:  [ -2, +2,  0, +3,  0,  0, +2, -3, -1, +1,  0,  0, -1 ]
  170918749:  [ +1, +4, -5, +1, +1, +1, +1,  0, -1, -1, +1,  0, -1 ]
  265326335:  [ -7, -2,  0, +1, -1, +1,  0, -2,  0, -1, +2, +1, +1 ]

  287080366:  [  0, +1, +1, -5, +2, +1,  0, -1, +3, -1, -1,  0,  0 ]
  362074049:  [ +5, -4, -2, +1,  0, -2,  0,  0, -2, +2, +2,  0,  0 ]
  587270881:  [ +4, +2, +1, +3, -1,  0, -1,  0,  0, +1, -1, -3, +1 ]

  617831551:  [ -6, +4, +2, +1,  0, -6,  0, +1,  0,  0, +1, +1,  0 ]
  740512499:  [ -1, -1, -5, -2, +7, -1,  0, +1,  0,  0, -1,  0,  0 ]
  831409151:  [ -9, -1, +2, -1, +3, +1,  0,  0, -1,  0, +2,  0, -2 ]

 1752438401:  [ +6, -4, +2,  0, -2, -2,  0, +2, -2,  0,  0, +1, +1 ]
 2151548801:  [ +6, -2, +2,  0,  0, -2,  0,  0, +2, -4, +1,  0, +1 ]
 2470954914:  [  0,  0, -1, +1, -3, -5, +2,  0,  0,  0, +3,  0, +1 ]
 3222617399:  [ -2, -6, -2, +4, +1, +2,  0, +2, -1,  0, -2,  0,  0 ]
```

Figure 31.4-B: Values $L(x)$ as linear combinations of logarithms of small primes.

Figure 31.4-B shows the linear combinations of logarithms of small primes that give the values $L(x)$. The first row gives the relation

$$L(51,744,295) \quad = \quad -2 \, \log(2) + 2 \, \log(3) + 3 \, \log(7) \pm \ldots - 1 \, \log(41) \tag{31.4-5}$$

The shown values, as a matrix, are the inverse the values in figure 31.4-A.

Precomputed logarithms of small primes can be used for the computation of the logarithms of integers k if one can determine a smooth number near k. For example, the logarithm of 65537 (a prime) can be computed as

$$\log(65537) \quad = \quad \log\left(65537 \cdot \frac{65536}{65536}\right) \; = \; \log\left(\frac{65537}{65536}\right) + \log(65536) \tag{31.4-6a}$$

$$= \quad \log\left(1 + \frac{1}{65536}\right) + 16 \, \log(2) \tag{31.4-6b}$$

The series of the first logarithm converges fast and $\log(2)$ is precomputed. This idea has been given by Jim White [priv.comm.]. If k is not near a smooth number but $u \cdot k$ is smooth where u factors into the chosen prime set, one can use the relation

$$\log(k) \quad = \quad \log(u\,k) - \log(u) \tag{31.4-7}$$

Here $\log(u)$ is the sum of precomputed logarithms and with $\log(u\,k)$ one can proceed as above.

Chapter 32

Numerical evaluation of power series

We give algorithms for the numerical evaluation of power series. When the series coefficients are rational the binary splitting (binsplit) algorithm can be applied for rational arguments, and the rectangular schemes with real (full-precision) arguments. As a special case of the binary splitting algorithm, a method for fast radix conversion is described. Finally we describe a technique for the summation of series with alternating coefficients.

32.1 The binary splitting algorithm for rational series

The straightforward computation of a series for which each term adds a constant amount of precision (for example, the arc-cotangent series with arguments > 1) to a precision of N digits involves the summation of proportional N terms. To get N bits of precision one has to add proportional N terms of the sum, each term involves one (length-N) short division (and one addition). Therefore the total work is proportional N^2, which makes it impossible to compute billions of digits from linearly convergent series even if they are as 'good' as Chudnovsky's famous series for π (given in [84]):

$$\frac{1}{\pi} = \frac{6541681608}{\sqrt{640320}^3} \sum_{k=0}^{\infty} \left(\frac{13591409}{545140134} + k \right) \left(\frac{(6k)!}{(k!)^3\,(3k)!} \frac{(-1)^k}{640320^{3k}} \right) \tag{32.1-1a}$$

$$= \frac{12}{\sqrt{640320}^3} \sum_{k=0}^{\infty} (-1)^k \frac{(6k)!}{(k!)^3\,(3k)!} \frac{13591409 + 545140134 \cdot k}{(640320)^{3k}} \tag{32.1-1b}$$

32.1.1 Binary splitting scheme for products

We motivate the binsplit algorithm by giving the analogue for the fast computation of the factorial. Define $f_{m,n} := m \cdot (m+1) \cdot (m+2) \cdots (n-1) \cdot n$, then $n! = f_{1,n}$. We compute $n!$ by recursively using the relation $f_{m,n} = f_{m,x} \cdot f_{x+1,n}$ where $x = \lfloor (m+n)/2 \rfloor$:

```
1   indent(i)=for(k=1,8*i,print1(" "));  \\ aux: print 8*i spaces
2
3   F(m, n, i=0)=
4   { /* Factorial, self documenting */
5       local(x, ret);
6       indent(i);  print( "F(", m, ", ", n, ")");
7       if ( m==n, /* then: */
8           ret = m;  \\ == F(m,m)
9       , /* else: */
10          x = floor( (m+n)/2 );
11          ret = F(m, x, i+1) * F(x+1, n, i+1);
12      );
```

```
F(1, 8)
        F(1, 4)
                F(1, 2)
                        F(1, 1)
                        ^== 1
                        F(2, 2)
                        ^== 2
                ^== 2
                F(3, 4)
                        F(3, 3)
                        ^== 3
                        F(4, 4)
                        ^== 4
                ^== 12
        ^== 24
        F(5, 8)
                F(5, 6)
                        F(5, 5)
                        ^== 5
                        F(6, 6)
                        ^== 6
                ^== 30
                F(7, 8)
                        F(7, 7)
                        ^== 7
                        F(8, 8)
                        ^== 8
                ^== 56
        ^== 1680
^== 40320
```

Figure 32.1-A: Quantities with the computation of 8!.

```
13      indent(i);  print( "^== ", ret);
14      return( ret );
15  }
```

The function prints the intermediate values occurring in the computation. The additional parameter `i` keeps track of the calling depth, used with the auxiliary function `indent()`. Figure 32.1-A shows the output with the computation of 8! =`F(1,8)`. A fragment like

```
F(5, 6)
        F(5, 5)
        ^== 5
        F(6, 6)
        ^== 6
^== 30
```

says "`F(5,6)` called `F(5,5)` [which returned 5], then called `F(6,6)` [which returned 6]. Then `F(5,6)` returned 30." For the computation of other products modify the line `ret=m;` as indicated in the code.

32.1.2 Binary splitting scheme for sums

For the evaluation of a sum $\sum_{k=0}^{N-1} a_k$ we use the ratios r_k of consecutive terms:

$$r_k \;:=\; \frac{a_k}{a_{k-1}} \tag{32.1-2}$$

Set $a_{-1} := 1$ to avoid a special case for $k = 0$. One has

$$\sum_{k=0}^{N-1} a_k \;=:\; r_0\,(1 + r_1\,(1 + r_2\,(1 + r_3\,(1 + \ldots (1 + r_{N-1})\ldots)))) \tag{32.1-3}$$

Now define

$$r_{m,n} \;:=\; r_m\,(1 + r_{m+1}\,(\ldots (1 + r_n)\ldots)) \quad \text{where} \quad m < n \tag{32.1-4a}$$

$$r_{m,m} \;:=\; r_m \tag{32.1-4b}$$

then

$$r_{m,n} \;=\; \frac{1}{a_{m-1}} \sum_{k=m}^{n} a_k \tag{32.1-5}$$

and especially

$$r_{0,n} \quad = \quad \sum_{k=0}^{n} a_k \tag{32.1-6}$$

We have

$$\begin{aligned} r_{m,n} &= r_m + r_m \cdot r_{m+1} + r_m \cdot r_{m+1} \cdot r_{m+2} + \ldots \qquad (32.1\text{-}7a) \\ &\quad \cdots + r_m \cdot \cdots \cdot r_x + r_m \cdot \cdots \cdot r_x \cdot [r_{x+1} + \cdots + r_{x+1} \cdot \cdots \cdot r_n] \\ &= r_{m,x} + \prod_{k=m}^{x} r_k \cdot r_{x+1,n} \qquad (32.1\text{-}7b) \end{aligned}$$

The product telescopes, one gets (for $m \le x < n$)

$$r_{m,n} \quad = \quad r_{m,x} + \frac{a_x}{a_{m-1}} \cdot r_{x+1,n} \tag{32.1-8}$$

32.1.3 Implementation using rationals

```
R(0, 6)
        R(0, 3)
                R(0, 1)
                        R(0, 0)
                        ^== 1/2
                        R(1, 1)
                        ^== 1/2
                ^== 3/4
                R(2, 3)
                        R(2, 2)
                        ^== 1/2
                        R(3, 3)
                        ^== 1/2
                ^== 3/4
        ^== 15/16
        R(4, 6)
                R(4, 5)
                        R(4, 4)
                        ^== 1/2
                        R(5, 5)
                        ^== 1/2
                ^== 3/4
                R(6, 6)
                ^== 1/2
        ^== 7/8
^== 127/128
```

Figure 32.1-B: Quantities with the binsplit computation of $\sum_{k=0}^{6} 2^{-(k+1)} = 127/128$.

Now we can formulate the binary splitting algorithm by giving a binsplit function using pari/gp:

```
1   R(m, n)=
2   { /* Rational binsplit */
3       local(x, ret);
4       if ( m==n, /* then: */
5           ret = A(m)/A(m-1);
6       , /* else: */
7           x = floor( (m+n)/2 );
8           ret = R(m, x) + A(x) / A(m-1) * R(x+1, n);
9       );
10      return( ret );
11  }
```

Here `A(k)` must be a function that returns the `k`-th term of the series we wish to compute, in addition one must have `a(-1)`=1. For example, to compute arctan(1/10) one would use

```
A(k)=if(k<0, 1, (-1)^(k)/((2*k+1)*10^(2*k+1)));
```

Figure 32.1-B shows the intermediate values with the computation of $\sum_{k=0}^{6} 2^{-(k+1)}$.

32.1.4 Implementation using integers

In case the programming language used does not provide rational numbers one needs to rewrite formula 32.1-8 in separate parts for denominator and numerator. With $a_i = p_i/q_i$, $p_{-1} = q_{-1} = 1$ and $r_{m,n} =: U_{m,n}/V_{m,n}$ one gets

$$U_{m,n} = p_{m-1}\, q_x\, U_{m,x}\, V_{x+1,n} + p_x\, q_{m-1}\, U_{x+1,n}\, V_{m,x} \tag{32.1-9a}$$

$$V_{m,n} = p_{m-1}\, q_x\, V_{m,x}\, V_{x+1,n} \tag{32.1-9b}$$

```
Q(0, 6)
        Q(0, 3)
                Q(0, 1)
                        Q(0, 0)
                        ^== [1, 10]
                        Q(1, 1)
                        ^== [-10, 3000]
                ^== [29900, 300000]
                Q(2, 3)
                        Q(2, 2)
                        ^== [3000, -500000]
                        Q(3, 3)
                        ^== [-500000, 70000000]
                ^== [-104250000000000000, 17500000000000000000]
        ^== [1569781275000000000000000000, 157500000000000000000000000000]
 [--snip--]
```

Figure 32.1-C: Explosive growth of intermediate quantities with computing arctan(1/10).

The following implementation also contains code for reduction to lowest terms:

```
1   Q(m, n)=
2   { /* Integer binsplit */
3       local(x, ret, bm, bx, tm, tx);
4       if ( m==n, /* then: */
5           bm = B(m);  bx = B(m-1);
6           ret = [ bm[1]*bx[2] , bx[1]*bm[2] ]; \\ == B(m)/B(m-1);
7           x = gcd(ret[1], ret[2]);  /* Reduction */
8           ret = [ret[1]/x, ret[2]/x];  /* Reduction */
9       , /* else: */
10          x = floor( (m+n)/2 );
11          tm = Q(m, x);     \\ [U_{m,x}, V_{m,x}]
12          tx = Q(x+1, n);   \\ [U_{x+1,n}, V_{m,n}]
13          bm = B(m-1);      \\ [p_{m-1}, q_{m-1}]
14          bx = B(x);        \\ [p_{x}, q_{x}]
15          \\ ret == Q(m, x) + B(x) / B(m-1) * Q(x+1, n);
16          ret = [ (bm[1]*bx[2]*tm[1]*tx[2] + bx[1]*bm[2]*tx[1]*tm[2])/10,
17                  (bm[1]*bx[2]*tm[2]*tx[2])/10 ];
18          x = gcd(ret[1], ret[2]);  /* Reduction */
19          ret = [ret[1]/x, ret[2]/x];  /* Reduction */
20      );
21      return( ret );
22  }
```

The reduction step can do good or bad, depending on the terms of the sum. When computing arctan(1/10) *without* the reduction, the intermediate quantities grow exponentially, as shown in figure 32.1-C. The square brackets are the quantities $[U_{m,n}, V_{m,n}]$. Such explosive growth will occur with all Taylor series unless the function argument equals one.

32.1.5 Performance

We compute the sum for arctan(1/10) up to the 5,000th term with the direct method, the rational binsplit and the integer binsplit with and without reduction. The timings for the computation are:

```
A(k)=if(k<0,1, (-1)^(k)/((2*k+1)*10^(2*k+1)));  \\ for rational binsplit
B(k)=if(k<0, [1,1], [(-1)^(k), ((2*k+1)*10^(2*k+1))] ); \\ for integer binsplit
N=5000;
```

```
  sum(k=0,N,A(k)); \\ direct method:  69,385 ms.
  R(0,N); \\ rational binsplit: 2,532 ms.
  Q(0,N); \\ integer binsplit with gcd reduction: 4,152 ms.
  Q(0,N); \\ integer binsplit without gcd reduction: >8min, "forever"
```

Things look quite different when computing the sum $\sum_{k=0}^{50,000} (-1)^k/(2k+1)^2$. The intermediate quantities U and V have only small common factors, so it is better to omit the reduction step:

```
B(k)=if(k<0, [1,1], [(-1)^k, (2*k+1)^2] );
A(k)=if(k<0,1, (-1)^(k)/(2*k+1)^2);
  N=50000;
  sum(k=0,N,A(k)); \\ direct method:  32,396 ms.
  R(0,N); \\ rational binsplit:  6,826 ms.
  Q(0,N); \\ integer binsplit with gcd reduction:  27,485 ms.
  Q(0,N); \\ integer binsplit without gcd reduction:  6,251 ms.
```

With builtin routines for binsplit summation the advantage will be much more in favor than these figures suggest.

The reason why summation via binary splitting is better than the straightforward way is that its complexity is only $O(\log N\, M(N))$, where $M(N)$ is the complexity of one N-bit multiplication (see [141]). If an FFT based multiplication algorithm is used ($M(N) = N\, \log N$) the work is $O((\log N)^2\, N)$ This means that sums of linear but sufficient convergence are again candidates for high precision computations. The algorithm should be implemented in the 'depth first' manner presented, and *not* via the naive pairs, pairs of pairs, etc. (breadth first) way. The reasons are better locality and less memory consumption. The naive way needs most memory after the first pass, when pairs have been multiplied.

32.1.6 Extending prior computations

The ratio $r_{0,N-1}$ (that is, the sum of the first N terms) can be reused if one wants to evaluate the sum to a higher precision than before. To get twice the precision use

$$r_{0,2\,N-1} \quad = \quad r_{0,N-1} + a_{N-1} \cdot r_{N,2\,N-1} \tag{32.1-10}$$

(this is formula 32.1-8 with $m = 0, x = N-1, n = 2N-1$). With explicit rational arithmetic:

$$U_{0,2N-1} \quad = \quad q_{N-1}\, U_{0,N-1}\, V_{N,2N-1} + p_{N-1}\, U_{N,2N-1}\, V_{0,N-1} \tag{32.1-11a}$$
$$V_{0,2N-1} \quad = \quad q_{N-1}\, V_{0,N-1}\, V_{N,2N-1} \tag{32.1-11b}$$

Thereby with the appearance of some new computer that can multiply two length $2 \cdot N$ numbers[1] one only needs to combine the two ratios $r_{0,N-1}$ and $r_{N,2N-1}$ that had been precomputed by the last generation of computers. This costs only a few full-size multiplications on your new and expensive supercomputer (instead of several *hundreds* for the iterative schemes), which means that one can improve on prior computations at low cost.

If one wants to stare at zillions of decimal digits of the floating point expansion then one division is also needed which costs no more than 4 multiplications as described in section 28.1 on page 559.

32.1.7 Computation of π: binary splitting versus AGM-type iterations

Using formula 32.1-1a on page 631 and the binary splitting scheme for the computation of π can outperform the AGM-style iterations given in section 30.4 on page 602. The reason is that the memory access pattern is more favorable than with the iterations. When computing N digits of π the iterations compute proportional $\log_2(N)$ roots (and or inverses) to full precision. At the last phase of each root computation full-length multiplications have to be computed. These access all memory storing a few full-precision words. In contrast, the binary splitting involves full-precision multiplications only at the very last phase.

[1] assuming the old model could multiply length-N numbers.

The drawback of the binary splitting scheme is that it may need significantly more memory than two full words. This may happen if the numerator and denominator grow fast which is more likely if no series so favorable as 32.1-1a can be used for the quantity to be computed. The problem can be mitigated (as pointed out by Richard Kreckel [priv.comm.], see [173], and [82]). by computing the floating point value whenever the integer values become too large. This technique is used in the CLN library [140].

32.1.8 Fast radix conversion

A binary splitting scheme for radix conversion of a radix-z integer $[a_N a_{N-1} \dots a_2 a_1 a_0]_z$ can be obtained via recursive application of the scheme

$$\sum_{k=M}^{N} a_k\, z^k \quad = \quad \sum_{k=M}^{M+X-1} a_k\, z^k + z^X \sum_{k=M+X}^{N} a_k\, z^k \tag{32.1-12}$$

where X is chosen to be the largest power of two that is smaller than $d := N - M$.

```
R(0, 15)
        R(0, 7)
                R(0, 3)
                        R(0, 1)
                        ^== 67
                        R(2, 3)
                        ^== 101
                ^== 25923
                R(4, 7)
                        R(4, 5)
                        ^== 7
                        R(6, 7)
                        ^== 33
                ^== 8455
        ^== 554132803
        R(8, 15)
                R(8, 11)
                        R(8, 9)
                        ^== 67
                        R(10, 11)
                        ^== 101
                ^== 25923
                R(12, 15)
                        R(12, 13)
                        ^== 7
                        R(14, 15)
                        ^== 33
                ^== 8455
        ^== 554132803
^== 2379982267079943491
```

Figure 32.1-D: Intermediate results when converting the number 2107654321076543_{16} to decimal.

We define an auxiliary function that computes (for $d > 1$) the largest exponent s so that $2^s < d$:

```
1   ex2le(d)=
2   { /* return largest s so that 2^s < d */
3       local(s, t);
4       t=1;  s=0;
5       while ( d>t,  t<<=1; s+=1; );
6       t >>= 1;  s--;
7       return(s);
8   }
```

We precompute z^2, z^4, z^8, …, z^{2^w} where $2^w < N$:

```
1   N=15;
2   z=16; \\ radix
3   vz=vector(ceil(log(N)/log(2)));
4   vz[1]=z;
5   for (k=2, length(vz), vz[k]=vz[k-1]^2);  \\ N space
```

Now the conversion function can be defined as

```
1   Ri(m, n, i=0)=
2   { /* Radix conversion, self documenting */
3       local(x, d, ret, t);
4       indent(i);  print( "R(", m, ", ", n, ")");
5       d = n-m;
6       if ( d <= 1, /* then: */
7           if ( d==0, ret = A(m); , ret = A(m) + z*A(n); );
8       , /* else: */
9           t = ex2le(d);
10          x = 1<<t;
11          ret = Ri(m, m+x-1, i+1) +  vz[t+1] * Ri(m+x, n, i+1);
12      );
13      indent(i);  print( "^== ", ret);
14      return( ret );
15  }
```

We define `A(k)=(k+3)%8;` and convert the 16-digit, radix-16 number $[a_{15}a_{14}\ldots a_2a_1a_0]_{16} := 2107654321076543_{16}$ to decimal. The intermediate results are shown in figure 32.1-D.

32.2 Rectangular schemes for evaluation of power series

The *rectangular scheme* for the evaluation of polynomials was given in [211], and later in [241]. We use it for the evaluation of truncated power series up to a given power $N-1$ of the series variable. We give two variants, one for series whose coefficients are small rationals (as for the logarithm), and another for series where the ratios of successive coefficients are a small rationals (as for the exponential function). When the numbers of rows and columns in the schemes are identical, a method involving proportional $\sqrt{N}$ full-precision multiplications is obtained. The schemes are very competitive up to very high precision in practice, even compared with AGM-based methods.

32.2.1 Rectangular scheme for arctan and logarithm

Computing the sum of the first N terms of a power series as

$$S_N \;:=\; \sum_{k=0}^{N-1} A_k\, z^k \;=\; A_0 + z\,(A_1 + z\,(A_2 + z\,(A_3 + \ldots z\,(A_{N-1})\ldots))) \tag{32.2-1}$$

costs N long (full-precision) multiplications if z is a full-precision number. If the A_k are small rational values, and $N = R\cdot C$ then one can rewrite S_N as

$$\begin{aligned} S_N \;=\;& A_{0C} + A_{0C+1}\,z + A_{0C+2}\,z^2 + \ldots + A_{1C-1}\,z^{C-1} + \\ &+z^C\,\left[A_{1C} + A_{1C+1}\,z + A_{1C+2}\,z^2 + \ldots + A_{2C-1}\,z^{C-1}+\right. \\ &+z^C\,\left[A_{2C} + A_{2C+1}\,z + A_{2C+2}\,z^2 + \ldots + A_{3C-1}\,z^{C-1}+\right. \\ &+z^C\,\left[A_{3C} + A_{3C+1}\,z + A_{3C+2}\,z^2 + \ldots + A_{4C-1}\,z^{C-1}+\right. \\ &+\quad \ldots \quad + \\ &+z^C\,\left[A_{(R-1)C} + A_{(R-1)C+1}\,z + A_{(R-1)C+2}\,z^2 + \ldots + A_{RC-1}\,z^{C-1}\right]\ldots]]] \end{aligned} \tag{32.2-2}$$

We compute S_N as

$$\left[\left[\left[\ldots[U_{R-1}]\;z^C + \ldots + U_3\right]\;z^C + U_2\right]\;z^C + U_1\right]\;z^C + U_0 \tag{32.2-3}$$

where $U_r := \sum_{k=0}^{C-1} A_{rC+k}\,z^k$ is the sum in one row of relation 32.2-2.

Precomputing the quantities z^2, z^3, z^4,$\ldots$ z^C involves $C-1$ long multiplications. The sums in each row of expression 32.2-2 involve only short multiplications with series coefficients A_i. The multiplication by z^C for each but the first row involves further $R-1$ long multiplications. The computation uses C temporaries (z, z^2, $\ldots$, z^C) and proportional $R+C$ long multiplications. Choosing $R = C = \sqrt{N}$ leads to a complexity of $2\sqrt{N}$ long multiplications, and also involves $\sqrt{N}$ temporaries.

32.2.1.1 Implementation for arctan

We implement the scheme for the arctan in pari/gp:

```
fa(n) =  \\ inverse of series coefficient
{  /* fa(n) := (-1)^n/(2n+1) */
    local(an);
    an = (2*n+1);
    if ( bitand(n,1), an=-an);
    return( an );
}
```

```
atan_rect(z, R, C)=
{  /* compute atan(z) as z*(1-z^2/3+z^4/5-z^6/7+-... +-z^(2*(R*C-1))/(2*R*C-1) */
    local(S, vz, s, ur, k);
    vz = vector(C);  \\ vz == [z^2,z^4,z^6,...,z^(2*C)]
    vz[1] = z*z;  \\ 1 long multiplication (special for arctan)
    for (k=2, C, vz[k]=vz[1]*vz[k-1]);  \\ C-1 long multiplications
    k = R*C;  \\ index of current coefficient
    s = 0;  \\ sum
    forstep (r=R-1, 0, -1,
        ur = 0;  \\ sum of this row
        forstep (c=C-1, 1, -1,  k-=1;  ur+=vz[c]/fa(k); );
        k -= 1;  ur += 1/fa(k);
        if ( r!=R-1, s*=vz[C]; );  \\ R-1 long multiplications
        s += ur;
    );
    s *= z;  \\ 1 long multiplication (special for arctan)
    return( s );
}
```

We compute $\pi/16$ as $\arctan(z)$ where $z = \sqrt{2\sqrt{2}+4} - \sqrt{2} - 1 \approx 0.19891236$ (using relation 31.1-30 on page 622 twice on $z = 1$), using a precision of 30,000 decimal digits. We use $R = C = \sqrt{N} =: S$:

```
? ? z=1;z=z+sqrt(z^2+1);z=z+sqrt(z^2+1);z=1/z; \\ ==> Pi/16
? a=atan(z);  \\ builtin arctan:  computed in 1,123 ms.
? r=atan_rect(z,S,S);  \\ computed in 2,377 ms.
  \\ using S=147, and N=S^2=21609
? a-r
0.E-30017  \\ result OK
```

The given implementation involves about two times of the cost of the builtin routine. Argument reduction make the method much more competitive:

```
? a=atan(z); \\ computed in 1,123 ms.
? z=1/z;
? for(k=1,32,z=z+sqrt(z^2+1)) \\ computed in 204 ms.
? z=1/z
4.57161899770987400328861548736 E-11
? S=ceil(sqrt(-1/2*rp*log(10)/log(z)))
39 \\ N=S^2=1521
? r=atan_rect(z,S,S); \\ computed in 284 ms.
? r*=2^32
? a-r
-1.3690050398194919519 E-30016  \\ OK
```

With 100,000 decimal digits the performance ratio is roughly the same. Note that one will have to limit the number C of temporaries according to the available memory.

Compute the inverse sine and cosine as $\arcsin(z) = \arctan \frac{z}{\sqrt{1-z^2}}$ and $\arccos(z) = \frac{\pi}{2} - \arcsin(z)$.

32.2.1.2 Implementation for the logarithm

A routine for $\log(1-z)$ is

```
log_rect(z, R, C)=
{ /* compute log(1-z) as 1+x/2+x^2/3+...+x^(R*C-1)/(R*C) */
    local(S, vz, s, ur, k);
    vz = vector(C);  \\ vz == [z^2,z^4,z^6,...,z^(2*C)]
    vz[1] = z;
    for (k=2, C, vz[k]=z*vz[k-1]);  \\ C-1 long multiplications
    k = R*C;  \\ index of current coefficient
```

```
 8       s = 0;  \\ sum
 9       forstep (r=R-1, 0, -1,
10           ur = 0;  \\ sum of this row
11           forstep (c=C-1, 1, -1,  k-=1;  ur+=vz[c]/(k+1); );
12           k -= 1;  ur += 1/(k+1);
13           if ( r!=R-1, s*=vz[C]; );  \\ R-1 long multiplications
14           s += ur;
15       );
16       s *= z;  \\ 1 long multiplication (special for arctan)
17       return( -s );
18   }
```

However, using a precision of 30,000 decimal digits and argument $z = 1/5$ the routine is slower than the builtin one (using the AGM) by a factor of about 1/7. With argument reduction (relation 31.1-27 on page 621), and $R = C = \sqrt{N} =: S$ we obtain a more competitive performance:

```
? e=log(1-z) \\ computed in 621 ms.
-0.22314355131420975576629509031 0
? z=1-z;
? for(k=1,32,z=sqrt(z));  \\ computed in 132 ms.
? z=1-z;  \\ == 5.19546566783481003872552738341 E-11
? S=ceil(sqrt(-rp*log(10)/log(z)))
55  \\ N=S^2=3025
? r=log_rect(z,N);  \\ computed in 461 ms.
? r*=2^32
-0.22314355131420975576629509031 0
? e-r
-6.071613129762050924 E-30008  \\ OK
```

We note that with both the logarithm and the arctan, subsequent computations with the builtin routine are faster as some constants that are computed with the first call are reused.

Compute the inverse hyperbolic sine and cosine as $\operatorname{arcsinh}(z) = \log(z + \sqrt{z^2+1})$ and $\operatorname{arccosh}(z) = \log(z + \sqrt{z^2-1})$ (for $z \geq 0$).

32.2.2 Rectangular scheme for exp, sine, and cosine

We rewrite the sum of the first N terms of a power series

$$S_N \quad := \quad \sum_{k=0}^{N-1} A_k \, z^k \tag{32.2-4}$$

as

$$\begin{aligned} S_N \quad = \quad & 1 \left[A_{0C+0} + A_{1C+0}\, z^{1C} + A_{2C+0}\, z^{2C} + \ldots + A_{(R-1)C+0} z^{(R-1)C} \right] + \\ & z^1 \left[A_{0C+1} + A_{1C+1}\, z^{1C} + A_{2C+1}\, z^{2C} + \ldots + A_{(R-1)C+1} z^{(R-1)C} \right] + \\ & z^2 \left[A_{0C+2} + A_{1C+2}\, z^{1C} + A_{2C+2}\, z^{2C} + \ldots + A_{(R-1)C+2} z^{(R-1)C} \right] + \\ & z^3 \left[A_{0C+3} + A_{1C+3}\, z^{1C} + A_{2C+3}\, z^{2C} + \ldots + A_{(R-1)C+3} z^{(R-1)C} \right] + \\ & + \quad \ldots \quad + \\ & z^{C-2} \left[A_{1C-2} + A_{2C-2}\, z^{1C} + A_{3C-2}\, z^{2C} + \ldots + A_{RC-2} z^{(R-1)C} \right] + \\ & z^{C-1} \left[A_{1C-1} + A_{2C-1}\, z^{1C} + A_{3C-1}\, z^{2C} + \ldots + A_{RC-1} z^{(R-1)C} \right] \end{aligned} \tag{32.2-5}$$

Compute the sum as (the transposed version of relation 32.2-2 on page 637)

$$S_N \quad = \quad [[[\ldots[[U_{C-1}]\, z + U_{C-2}]\, z + \ldots + U_3]\, z + U_2]\, z + U_1]\, z + U_0 \tag{32.2-6}$$

where $U_c = \sum_{k=0}^{R-1} A_{kC+c}\, z^{kC}$ (C temporary sums are computed). When proceeding colum-wise the update $A_i \to A_{i+1}$ involves only a short multiplication by the ratio A_{i+1}/A_i. Only when going to the next column a long multiplication by z^C is required ($R-1$ long multiplications). Finally, there are $C-1$ long multiplications by z.

32.2.2.1 Implementation for the exponential function

A routine for the computation of $\exp(z) - 1$ can be given as follows:

```
exp_rect(z, R, C)=
{ /* compute exp(z)-1 as z*[ 1+z/2!+z^2/3! +...+z^(R*C-1)/((R*C)!) ] */
    local(ur, zc, k, t);
    zc = z^C; \\ proportional log(C) long multiplications
    ur = vector(C);
    k = 1;  \\ ratio of series coefficients  /* set to zero for plain exp */
    t = 1.0;
    for (r=1, R,  \\ number of columns (!)
        for (c=1, C,  ur[c] += t;  k++;  t /= (k); );
        if ( r!=R,  t *= zc; );  \\ R-1 long multiplications
    );
    t = ur[C];
    forstep (c=C-1, 1, -1,  t*=z; t+=ur[c]);  \\ C-1 long multiplications
    t *= z;  /* omit for plain exp */
    return( t );
}
```

We use the argument reduction given as relation 31.2-18 on page 625 and compute $exp(1/5)$ to a precision of 30,000 decimal digits. We use $R = C = \sqrt{N} =: S$:

```
? z=0.2;
? e=exp(z)  \\ computed in 855 ms.
1.22140275816016983392107199464
? nred=32;
? z/=2^nred
4.65661287307739257812500000000 E-11
? S=48; \\ N=S^2=2304
? r=exp_rect(z,S,S)  \\ computed in 395 ms.
4.65661287318581279537523334667 E-11
? for(k=1,nred,r=r+r+r^2);  \\ computed in 68 ms.
? r+=1
1.22140275816016983392107199464
? e-r
7.965120231677044083 E-30016  \\ OK
```

Using 100,000 digits, `nred=112`, and `S=52` we obtain the timings

```
? e=exp(z);                \\ computed in 8,601 ms.
? r=exp_rect(z,S,S);       \\ computed in 2,345 ms.
? for(k=1,nred,r=r+r+r^2); \\ computed in 1,640 ms.
```

32.2.2.2 Implementation for the cosine

A routine for computing $\cos(z) - 1$ can be given as

```
cos_rect(z, R, C)=
{ /* compute cos(z)-1 as z^2*[ -1/2!+z^2/4! - z^4/6! +- ... ] */
    local(ur, zc, k, t);
    z *= z;
    zc = z^C; \\ proportional log(C) long multiplications
    ur = vector(C);
    k = 2;  \\ ratio of series coefficients
    t = -0.5;
    for (r=1, R,  \\ number of columns (!)
        for (c=1, C,  ur[c] += t;  k++;  t /= (k); k++;  t /= -(k); );
        if ( r!=R,  t *= zc; );  \\ R-1 long multiplications
    );
    t = ur[C];
    forstep (c=C-1, 1, -1,  t*=z; t+=ur[c]);  \\ C-1 long multiplications
    t *= z;  /* omit for plain exp */
    return( t );
}
```

We use the argument reduction as in relation 31.2-20 on page 625, and compute $\cos(1/5)$ to 30,000 decimal digits:

```
? z=0.2;
? e=cos(z)  \\ computed in 788 ms.
0.98006657784124163112419651674 8
? nred=32;
```

```
? z/=2^nred;
? S=34;  \\ N=S^2=1156
? r=cos_rect(z,S,S);  \\ computed in 318 ms.
? for(k=1,nred,r=2*(r+1)^2-2);  \\ computed in 70 ms.
? r+=1
0.98006657784124163112419651674 8
? e-r
-3.646143951667310362 E-30017  \\ OK
```

The sine and tangent can be computed as $\sin(z) = \sqrt{1-\cos(z)^2}$, and $\tan(z) = \sin(z)/\cos(z)$.

The routine is easily converted to compute the hyperbolic cosine. The relation $\exp(z) = \cosh(z) - \sqrt{\cosh(z)^2 - 1}$ gives an alternative way to compute the exponential function.

32.3 The magic sumalt algorithm for alternating series

The following convergence acceleration algorithm for alternating series is due to Cohen, Villegas and Zagier, see [92]. As remarked in the cited paper, the algorithm often gives meaningful results also for non-alternating and even divergent series.

The algorithm computes an estimate of the sum $s = \sum_{k=0}^{\infty} x_k$ as

$$s_n \quad = \quad \sum_{k=0}^{n-1} c_{n,k}\, x_k \tag{32.3-1}$$

The weights $c_{n,k}$ do not depend on the values x_j. With the following pseudo code the summands x_k have to be supplied in the array `x[0,1,...,n-1]`:

```
1   function sumalt(x[], n)
2   {
3       d := (3+sqrt(8))^n
4       d := (d+1/d)/2
5       b := 1
6       c := d
7       s := 0
8       for k:=0 to n-1
9       {
10          c := c - b
11          s := s + c * x[k]
12          b := b * (2*(n+k)*(n-k)) / ((2*k+1)*(k+1))
13      }
14      return s/d
15  }
```

With alternating sums the accuracy of the estimate will be $(3+\sqrt{8})^{-n} \approx 5.82^{-n}$. For example, the estimate for $4 \cdot \arctan(1)$ using the first 8 terms is

$$\pi \quad \approx \quad 4 \cdot \left(\frac{1}{1} - \frac{1}{3} + \frac{1}{5} - \frac{1}{7} + \frac{1}{9} - \frac{1}{11} + \frac{1}{13} - \frac{1}{15} \right) \; = \; 3.017\ldots \tag{32.3-2}$$

The sumalt-massaged estimate is

$$\begin{aligned} \pi \quad &\approx \quad 4 \cdot \left(\frac{665856}{1} - \frac{665728}{3} + \frac{663040}{5} - \frac{641536}{7} + \right. \\ &\qquad \left. + \frac{557056}{9} - \frac{376832}{11} + \frac{163840}{13} - \frac{32768}{15} \right) / 665857 \\ &= \quad 4 \cdot 3365266048/4284789795 = 3.141592665\ldots \end{aligned} \tag{32.3-3}$$

and already gives 7 correct digits of π. The linear but impressive growth of the accuracy of successive sumalt estimates with n, the number of terms used, is illustrated in figure 32.3-A.

n:	sumalt(n)	sumalt(n)$-\pi$
1:	2.66666666666666666666	0.474925986923126571795
2:	3.13725490196078431372	0.004337751629008924737
3:	3.14074074074074074074	0.000851912849052497721
4:	3.14163571841214822150 7	−0.000043064822354983044
5:	3.14158654640367396834 8	0.000006107186119270114
6:	3.14159344421566940366 0	−0.000000690625866165197
7:	3.14159256493754012201 5	0.000000088652253116447
8:	3.14159266522431586401 7	−0.000000011634522625555
9:	3.14159265200881195161 9	0.000000001580981286843
10:	3.14159265380973156931 8	−0.000000000219938330856
11:	3.14159265355857875551 3	0.000000000031214482948
12:	3.14159265359429633847 0	−0.000000000004503100007
13:	3.14159265358913458051 7	0.000000000000658657944
14:	3.14159265358989071862 5	−0.000000000000097480163
15:	3.14159265358977866437 5	0.000000000000014574087
16:	3.14159265358979543677 5	−0.000000000000002198312
17:	3.14159265358979290428 5	0.000000000000000334177
18:	3.14159265358979328961 4	−0.000000000000000051151
19:	3.14159265358979323058 4	0.000000000000000007877
20:	3.14159265358979323968 2	−0.000000000000000001220

Figure 32.3-A: Sumalt-estimates of $\pi = 4 \cdot \arctan(1)$ using $n = 1, 2, \ldots, 20$ terms.

Therefore even slowly converging series like

$$\pi = 4 \cdot \sum_{k=0}^{\infty} \frac{(-1)^k}{2\,k+1} = 4 \cdot \arctan(1) \tag{32.3-4a}$$

$$C = \sum_{k=0}^{\infty} \frac{(-1)^k}{(2\,k+1)^2} = 0.9159655941772190\ldots \tag{32.3-4b}$$

$$\log(2) = \sum_{k=0}^{\infty} \frac{(-1)^k}{k+1} = 0.6931471805599453\ldots \tag{32.3-4c}$$

$$\zeta(s) = \frac{1}{1-2^{1-s}} \sum_{k=1}^{\infty} \frac{(-1)^k}{k^s} \tag{32.3-4d}$$

can be used to compute estimates that are correct up to thousands of digits. The algorithm scales like n^2 if the series terms in the array `x[]` are small rational values and like $n^3 \cdot \log(n)$ if they are full precision (rational or float) values.

In fact, pari/gp has a built-in sumalt routine, we use it to compute the *Catalan constant*:

```
? default(realprecision,1000);
? sumalt(k=0,(-1)^k/(2*k+1)^2);  \\ takes    60 ms.
? default(realprecision,2000);
? sumalt(k=0,(-1)^k/(2*k+1)^2);  \\ takes   376 ms.
? default(realprecision,4000);
? sumalt(k=0,(-1)^k/(2*k+1)^2);  \\ takes 2,730 ms.
```

The time scales roughly with the third power of the precision used.

All values c_k and b_k occurring in the computation are integers. In fact, the b_k in the computation with n terms are the coefficients of the expanded n-th Chebyshev polynomial (of the first kind, see section 34.2 on page 667) with argument $1 + 2x$:

k :	b_k	c_k
0:	1	665857
1:	128	665856
2:	2688	665728
3:	21504	663040
4:	84480	641536
5:	180224	557056
6:	212992	376832
7:	131072	163840
8:	32768	32768

$$\begin{aligned} T_8(1+2x) &= 1 + 128x + 2688x^2 + 21504x^3 + 84480x^4 + \qquad (32.3\text{-}5a)\\ &\quad +180224x^5 + 212992x^6 + 131072x^7 + 32768x^8 = T_{16}(\sqrt{1+x}) \\ T_{16}(x) &= 1 - 128x^2 + 2688x^4 - 21504x^6 + 84480x^8 - \qquad (32.3\text{-}5b)\\ &\quad -180224x^{10} + 212992x^{12} - 131072x^{14} + 32768x^{16} \end{aligned}$$

Now observe that one has always $c_n = b_n = 2^{2n-1}$ in a length-n sumalt computation. Obviously, 'going backwards' avoids the computation of $(3+\sqrt{8})^n$:

```
1   function sumalt(x[], n)
2   {
3       b := 2**(2*n-1)
4       c := b
5       s := 0
6       for k:=n-1 to 0 step -1
7       {
8           s := s + c * x[k]
9           b := b * ((2*k+1)*(k+1)) / (2*(n+k)*(n-k))
10          c := c + b
11      }
12      return s/c
13  }
```

The b_k and c_k occurring in a length-n sumalt computation can be given explicitly as

$$b_k = \frac{n}{n+k}\binom{n+k}{2k}2^{2k} \qquad (32.3\text{-}6a)$$

$$c_k = \sum_{i=k}^{n}\frac{n}{n+i}\binom{n+i}{2i}2^{2i} \qquad (32.3\text{-}6b)$$

To compute an estimate of $\sum_{k=0}^{\infty} x_k$ using the first n partial sums use the following pseudo code (the partial sums $p_k = \sum_{j=0}^{k} x_j$ are expected in `p[0,1,...,n-1]`):

```
1   function sumalt_partial(p[], n)
2   {
3       d := (3+sqrt(8))^n
4       d := (d+1/d)/2
5       b := 1
6       c := d
7       s := 0
8       for k:=0 to n-1
9       {
10          s := s + b * p[k]
11          b := b * (2*(n+k)*(n-k)) / ((2*k+1)*(k+1))
12      }
13      return s/d
14  }
```

The backward variant is:

```
1   function sumalt_partial(p[], n)
2   {
3       b := 2**(2*n-1)
4       c := b
5       s := 0
6       for k:=n-1 to 0 step -1
```

```
7       {
8           s := s + b * p[k]
9           b := b * ((2*k+1)*(k+1)) / (2*(n+k)*(n-k))
10          c := c + b
11      }
12      return s/c
13  }
```

For series of already geometrical rate of convergence (where $|a_k/a_{k+1}| \approx e$) it is better to use

```
1   function sumalt_partial(p[], n, e)
2   {
3       d := ( 2*e + 1 + 2*sqrt(e*(e+1)) )^n
4       d := (d+1/d)/2
5       b := 1
6       c := d
7       s := 0
8       for k:=0 to n-1
9       {
10          s := s + b * p[k]
11          b := b * (2*(n+k)*(n-k)) / ((2*k+1)*(k+1)) * e
12      }
13      return s/d
14  }
```

Convergence is improved from $\sim e^{-n}$ to $\sim (2e + 1 + 2\sqrt{e(e+1)})^{-n} \approx (4e + 2)^{-n}$. This algorithm specializes to the original one for $e = 1$.

Implementations of the sumalt algorithm and the variant for partial sums are given in [hfloat: src/hf/sumalt.cc].

Chapter 33

Computing the elementary functions with limited resources

This chapter presents two types of algorithms for computations with limited resources, the shift-and-add and the CORDIC algorithms. The algorithms allow the computations of the elementary functions as the logarithm, exponential function, sine, cosine and their inverses with only shifts, adds, comparisons and table lookups. Algorithms of this type are usually used for pocket calculators.

33.1 Shift-and-add algorithms for $\log_b(x)$ and b^x

In this section so-called *shift-and-add algorithms* for the computation of $\log_b(x)$ and b^x are presented. These algorithms use only additions, multiplications by a power of two ('shifts') and comparisons. Precomputed lookup table with as many entries as the desired accuracy in bits is required. The algorithms are especially useful with limited hardware capabilities.

The implementations given in this section use floating point numbers. They can be rewritten to scaled integer arithmetic without difficultly.

33.1.1 Computing the base-b logarithm

We use a table that contains the values $A_n = \log_b\left(1 + \frac{1}{2^n}\right)$ where $n > 0$, it is created as follows:

```
1   double *shiftadd_ltab;
2   ulong ltab_n;
3
4   void
5   make_shiftadd_ltab(double b)
6   {
7       double l1b = 1.0 / log(b);
8       double s = 1.0;
9       for (ulong k=0; k<ltab_n; ++k)
10      {
11          shiftadd_ltab[k] = log(1.0+s) * l1b;  // == log_b(1+1/2^k)
12          s *= 0.5;
13      }
14  }
```

The algorithm takes as input the argument $x \geq 1$ and the number of iterations N and computes $\log_b(x)$. It proceeds as follows:

1. Initialize: set $t_0 = 0$, $e_0 = 1$.

n :	u_n	t_n	e_n	A_n
init	-	0.00000000	+1.00000000	+1.00000000
1:	1.50000000	0.00000000	+1.00000000	+0.58496250
2:	1.25000000	0.00000000	+1.00000000	+0.32192809
2:	1.56250000	0.32192809	+1.25000000	+0.32192809
3:	1.40625000	0.32192809	+1.25000000	+0.16992500
3:	1.58203125	0.49185309	+1.40625000	+0.16992500
4:	1.49414062	0.49185309	+1.40625000	+0.08746284
5:	1.45019531	0.49185309	+1.40625000	+0.04439411
6:	1.42822265	0.49185309	+1.40625000	+0.02236781
7:	1.41723632	0.49185309	+1.40625000	+0.01122725
8:	1.41174316	0.49185309	+1.40625000	+0.00562454
8:	1.41725778	0.49747764	+1.41174316	+0.00562454
9:	1.41450047	0.49747764	+1.41174316	+0.00281501
10:	1.41312181	0.49747764	+1.41174316	+0.00140819
10:	1.41450182	0.49888583	+1.41312181	+0.00140819
11:	1.41381182	0.49888583	+1.41312181	+0.00070426
11:	1.41450215	0.49959010	+1.41381182	+0.00070426
12:	1.41415698	0.49959010	+1.41381182	+0.00035217
12:	1.41450224	0.49994228	+1.41415698	+0.00035217
13:	1.41432961	0.49994228	+1.41415698	+0.00017609
14:	1.41424330	0.49994228	+1.41415698	+0.00008805
15:	1.41420014	0.49994228	+1.41415698	+0.00004402
15:	1.41424330	0.49998631	+1.41420014	+0.00004402
∞:	1.41421356	0.50000000	+1.41421356	+0.00000000
	$= x$	$= \log_2(\sqrt{2})$	$= x$	$= 0$

Figure 33.1-A: Numerical values occurring in the shift-and-add computation of $\log_2(\sqrt{2}) = 1/2$. The computation of $\log_{1/2}(\sqrt{2}) = -1/2$ corresponds to the same values but opposite signs for all entries A_n and y_n.

2. Compute $u_n = e_n \cdot (1 + 2^{-n})$. If $u_n \leq x$ the set $d_n = 1$, else set $d_n = 0$.

3. If $d_n \neq 0$ then set $t_{n+1} = t_n + A_n$ and $e_{n+1} = u_n$ and repeat the last step. Else set $t_{n+1} = t_n$ and $e_{n+1} = e_n$.

4. Increment n. If $n = N$ return t_n, else goto step 2.

A C++ implementation is given in [FXT: arith/shiftadd-log-demo.cc], note that the variable `n` equals N, and `k` equals n:

```
double
shiftadd_log(double x, ulong n)
{
    if ( n>=ltab_n )  n = ltab_n;
    double t = 0.0;
    double e = 1.0;
    double v = 1.0;
    // [PRINT]
    for (ulong k=1; k<n; ++k)
    {
        v *= 0.5;  // v == (1>>k)

        double u;
        bool d;
        while ( 1 )
        {
            u = e + e * v;  // u=e;  u+=(e>>k);
            d = ( u<=x );
            // [PRINT]
            if ( d==false )  break;
            t += shiftadd_ltab[k];
            e = u;
        }
    }
```

```
25        return  t;
26    }
```

The variable `v` is a power of 1/2 so all multiplies by it can with scaled integer arithmetic be replaced by shifts as indicated by the comments. The values for first steps of the computation for the argument $x_0 = \sqrt{2}$ are given in figure 33.1-A. The columns of the figure correspond to the variables `k`($= n$), `u`($= u_n$), `t`($= t_n$), `e`($= e_n$), and `shiftadd_ltab[k]`($= A_n$).

n :	u_n	t_n	e_n	A_n
init	-	0.00000000	+1.00000000	+1.00000000
1:	1.50000000	0.00000000	+1.00000000	+0.58496250
1:	2.25000000	0.58496250	+1.50000000	+0.58496250
1:	3.37500000	1.16992500	+2.25000000	+0.58496250
1:	5.06250000	1.75488750	+3.37500000	+0.58496250
1:	7.59375000	2.33985000	+5.06250000	+0.58496250
1:	11.3906250	2.92481250	+7.59375000	+0.58496250
2:	9.49218750	2.92481250	+7.59375000	+0.32192809
3:	8.54296875	2.92481250	+7.59375000	+0.16992500
4:	8.06835937	2.92481250	+7.59375000	+0.08746284
5:	7.83105468	2.92481250	+7.59375000	+0.04439411
5:	8.07577514	2.96920662	+7.83105468	+0.04439411
6:	7.95341491	2.96920662	+7.83105468	+0.02236781
6:	8.07768702	2.99157443	+7.95341491	+0.02236781
∞:	8.00000000	2.99999999	+8.00000000	+0.00000000
	$= x$	$= \log_2(8)$	$= x$	$= 0$

Figure 33.1-B: Values occurring in the first few steps of a shift-and-add computation of $\log_2(8) = 3$.

The algorithm has been adapted from [205] (chapter 5) where the correction is made only once for each value A_n limiting the range of convergence to $x < X$ where

$$\begin{aligned} X &= \prod_{n=0}^{\infty} \left(1 + \frac{1}{2^n}\right) \\ &= 4.76846205806274344829979857735679447754 3\ldots \end{aligned} \tag{33.1-1}$$

As given, the algorithm converges for any $x > 0$, $x \neq 1$. A numerical example for the argument $x = 8$ is given in figure 33.1-B. The basis b must satisfy $b > 0$ and $b \neq 1$.

33.1.2 Computing b^x

We can use the same precomputed table as with the computation of $\log_b(x)$.

The algorithm takes as input the argument x and the number of iterations N and computes b^x for $b > 1$, $x \in \mathbb{R}$. It proceeds as follows:

1. Initialize: set $t_0 = 0$, $e_0 = 1$.
2. Compute $u_n = t_n + A_n$. If $u_n \leq x$ the set $d_n = 1$, else set $d_n = 0$.
3. If $d_n \neq 0$ then set $t_{n+1} = u_n$ and $e_{n+1} = e_n \cdot (1 + 2^{-n})$ and repeat the last step. Else set $t_{n+1} = t_n$ and $e_{n+1} = e_n$.
4. Increment n. If $n = N$ return e_n, else goto step 2.

A C++ implementation is given in [FXT: arith/shiftadd-exp-demo.cc]:

```
1   double
2   shiftadd_exp(double x, ulong n)
3   {
4       if ( n>=ltab_n )  n = ltab_n;
```

n :	u_n	t_n	e_n	A_n
init	0.00000000	0.00000000	+1.00000000	+0.00000000
1:	0.58496250	0.00000000	+1.00000000	+0.58496250
2:	0.32192809	0.00000000	+1.00000000	+0.32192809
2:	0.64385618	0.32192809	+1.25000000	+0.32192809
3:	0.49185309	0.32192809	+1.25000000	+0.16992500
3:	0.66177809	0.49185309	+1.40625000	+0.16992500
4:	0.57931593	0.49185309	+1.40625000	+0.08746284
5:	0.53624721	0.49185309	+1.40625000	+0.04439411
6:	0.51422090	0.49185309	+1.40625000	+0.02236781
7:	0.50308035	0.49185309	+1.40625000	+0.01122725
8:	0.49747764	0.49185309	+1.40625000	+0.00562454
8:	0.50310219	0.49747764	+1.41174316	+0.00562454
9:	0.50029266	0.49747764	+1.41174316	+0.00281501
10:	0.49888583	0.49747764	+1.41174316	+0.00140819
10:	0.50029403	0.49888583	+1.41312181	+0.00140819
11:	0.49959010	0.49888583	+1.41312181	+0.00070426
11:	0.50029437	0.49959010	+1.41381182	+0.00070426
12:	0.49994228	0.49959010	+1.41381182	+0.00035217
12:	0.50029446	0.49994228	+1.41415698	+0.00035217
13:	0.50011838	0.49994228	+1.41415698	+0.00017609
14:	0.50003033	0.49994228	+1.41415698	+0.00008805
15:	0.49998631	0.49994228	+1.41415698	+0.00004402
15:	0.50003034	0.49998631	+1.41420014	+0.00004402
∞:	0.50000000	0.50000000	+1.41421356	+0.00000000
	$= x$	$= x$	$= 2^{1/2}$	$= 0$

Figure 33.1-C: Numerical values occurring in the shift-and-add computation of $b^x = 2^{1/2} = \sqrt{2}$. The values are printed at points where a comment `[PRINT]` appears in the code.

```
5       double t = 0.0;
6       double e = 1.0;
7       double v = 1.0;
8       // [PRINT]
9       for (ulong k=1; k<n; ++k)
10      {
11          v *= 0.5;  // v == (1>>k)
12
13          double u;
14          bool d;
15          while ( 1 )
16          {
17              u = t + shiftadd_ltab[k];
18              d = ( u<=x );
19              // [PRINT]
20              if ( d==false )  break;
21              t = u;
22              e += e * v;  // e+=(e>>k);
23          }
24      }
25      return  e;
26  }
```

33.1.3 An alternative algorithm for the logarithm

A slightly different method for the computation of the base-b logarithm ($b > 0$, $b \neq 1$) is given in [167]. Here the table used has to contain the values $A_n = \log_b\left(\frac{2^n}{2^n-1}\right)$ where $n > 0$:

```
1   double *briggs_ltab;
2   ulong  ltab_len;
3
4   void
5   make_briggs_ltab(ulong na, double b)
```

n :	x_n	y_n	z_n	A_n
init	1.41421356	0.00000000	+0.70710678	+0.00000000
2:	1.41421356	0.00000000	+0.35355339	+0.41503749
2:	1.06066017	0.41503749	+0.26516504	+0.41503749
3:	1.06066017	0.41503749	+0.13258252	+0.19264507
4:	1.06066017	0.41503749	+0.06629126	+0.09310940
5:	1.06066017	0.41503749	+0.03314563	+0.04580368
5:	1.02751454	0.46084118	+0.03210982	+0.04580368
6:	1.02751454	0.46084118	+0.01605491	+0.02272007
6:	1.01145962	0.48356126	+0.01580405	+0.02272007
7:	1.01145962	0.48356126	+0.00790202	+0.01131531
7:	1.00355759	0.49487657	+0.00784029	+0.01131531
8:	1.00355759	0.49487657	+0.00392014	+0.00564656
9:	1.00355759	0.49487657	+0.00196007	+0.00282051
9:	1.00159752	0.49769709	+0.00195624	+0.00282051
10:	1.00159752	0.49769709	+0.00097812	+0.00140957
10:	1.00061940	0.49910666	+0.00097716	+0.00140957
11:	1.00061940	0.49910666	+0.00048858	+0.00070461
11:	1.00013081	0.49981128	+0.00048834	+0.00070461
12:	1.00013081	0.49981128	+0.00024417	+0.00035226
13:	1.00013081	0.49981128	+0.00012208	+0.00017612
13:	1.00000873	0.49998740	+0.00012207	+0.00017612
∞:	1.00000000	0.50000000	+0.00000000	+0.00000000
	$= 1$	$= \log_2(\sqrt{2})$	$= 0$	$= 0$

Figure 33.1-D: Numerical values occurring in the computation of $\log_2(\sqrt{2}) = 1/2$. The value of n is incremented in the inner loop (comment `[PRINT1]` in the code, the value of z changes). The values of x and y change just before the location of the comment `[PRINT2]`, corresponding to consecutive rows with same value of n. The computation of $\log_{1/2}(\sqrt{2}) = -1/2$ corresponds to the same values but opposite signs for all entries A_n and y_n.

```
6   {
7       double l1b = 1.0 / log(b);
8       double s = 2.0; // == 2^k
9       briggs_ltab[0] = -1.0; // unused
10      for (ulong k=1; k<na; ++k)
11      {
12          briggs_ltab[k] = log(s/(s-1.0)) * l1b;
13          s *= 2.0;
14      }
15  }
```

The algorithm terminates when a given precision (`eps`) is reached:

```
1   double
2   briggs_log(double x, double eps)
3   {
4       double y = 0;
5       double z = x * 0.5;
6       // [PRINT]
7
8       ulong k = 1;
9       double v = 0.5;  // v == 2^(-k)
10      while ( fabs(x-1.0)>=eps )
11      {
12          while ( fabs(x-z)<1.0 )
13          {
14              z *= 0.5;
15              ++k;  v *= 0.5;
16              if ( k >= ltab_len )  goto done;  // no more table entries
17              // [PRINT1]
18          }
19
20          x -= z;
21          y += briggs_ltab[k];
22          z = x * v;  // z=(x>>k)
```

```
23          // invariant:  y_k + log_b(x_k) == log_b(x_0)
24
25          // [PRINT2]
26      }
27
28    done:
29      return  y;
30  }
```

The code is given in [FXT: arith/briggs-log-demo.cc]. The values for first steps of the computation for the argument $x_0 = \sqrt{2}$ are given in figure 33.1-D. The argument x must be greater than or equal to 1. Knuth [167] gives $1 \le x < 2$ but the restriction to values smaller than 2 does not seem to be necessary.

33.2 CORDIC algorithms

The so-called CORDIC algorithms can be used for the computation of functions like sine, cosine, exp and log. The acronym CORDIC stands for **Co**ordinate **R**otation **Di**gital **C**omputer.

Similar to the shift-and-add algorithms (section 33.1) only multiplications by powers of two (shifts), additions, subtractions and comparisons are used. Again, a precomputed lookup table with as many entries as the desired accuracy in bits is required.

Some early floating point units (FPUs) used CORDIC algorithms and your pocket calculator surely does.

33.2.1 The circular case: sine and cosine

n :	x_n	y_n	z_n	$-d \cdot A_n$
init	0.60725293	0.00000000	+1.04719755	+0.00000000
0:	0.60725293	0.60725293	+0.26179938	-0.78539816
1:	0.30362646	0.91087940	-0.20184822	-0.46364760
2:	0.53134631	0.83497278	+0.04313044	+0.24497866
3:	0.42697471	0.90139107	-0.08122455	-0.12435499
4:	0.48331166	0.87470515	-0.01880574	+0.06241880
5:	0.51064619	0.85960166	+0.01243409	+0.03123983
6:	0.49721492	0.86758051	-0.00318963	-0.01562372
7:	0.50399289	0.86369602	+0.00462270	+0.00781234
8:	0.50061908	0.86566474	+0.00071647	-0.00390623
9:	0.49892833	0.86664251	-0.00123664	-0.00195312
10:	0.49977466	0.86615528	-0.00026008	+0.00097656
11:	0.50019758	0.86591124	+0.00022819	+0.00048828
12:	0.49998618	0.86603336	-0.00001594	-0.00024414
13:	0.50009190	0.86597233	+0.00010612	+0.00012207
14:	0.50003904	0.86600285	+0.00004508	-0.00006103
15:	0.50001261	0.86601811	+0.00001457	-0.00003051
∞:	0.50000000	0.86602540	+0.00000000	+0.00000000
	$= \cos(\pi/3)$	$= \sin(\pi/3)$	$= 0$	$= 0$

Figure 33.2-A: Numerical values occurring in the CORDIC computation of $\cos(\pi/3)$ and $\sin(\pi/3)$.

We start with a CORDIC routine for the computation of the sine and cosine. The lookup table has to contain the values $\arctan(2^{-n})$ for $n = 0, 1, 2, 3, \ldots$, these shall be stored in the array `cordic_ctab[]`. An implementation of the function is given in [FXT: arith/cordic-circ-demo.cc]:

```
1  void
2  cordic_circ(double theta, double &s, double &c, ulong n)
3  {
4      double x = cordic_1K;
5      double y = 0;
6      double z = theta;
7      double v = 1.0;
```

```
 8       // [PRINT]
 9       for (ulong k=0; k<n; ++k)
10       {
11           double d = ( z>=0 ? +1 : -1 );
12           double tx = x - d * v * y;
13           double ty = y + d * v * x;
14           double tz = z - d * cordic_ctab[k];
15           x = tx;  y = ty;  z = tz;
16           v *= 0.5;
17           // [PRINT]
18       }
19       c = x;
20       s = y;
21   }
```

For the sake of clarity floating point types are used. All operations can easily be converted to integer arithmetic. The multiplications by `d` are sign changes and should be replaced by an `if`-construct. The multiplications by `v` are shifts.

The values for first 16 steps of the computation for the argument $z_0 = \theta = \pi/3 = 1.04719755\ldots$ are given in figure 33.2-A. While z gets closer to zero (however, the magnitude of z does not necessarily decrease with every step) the values of x and y approach $\sin(\pi/3) = 1/2$ and $\cos(\pi/3) = \sqrt{3}/2 = 0.86602540\ldots$, respectively.

More formally, one initializes

$$x_0 = 1/K = 0.607252935008881\ldots \tag{33.2-1a}$$
$$y_0 = 0 \tag{33.2-1b}$$
$$z_0 = \theta \tag{33.2-1c}$$

and iterates (starting with $n = 0$)

$$A_n = \arctan\left(2^{-n}\right) \quad \text{(precomputed)} \tag{33.2-1d}$$
$$v_n = 2^{-n} \tag{33.2-1e}$$
$$d_n = \operatorname{sign}(z_n) \tag{33.2-1f}$$
$$x_{n+1} = x_n - d_n\, v_n\, y_n \quad \to \cos(\theta) \tag{33.2-1g}$$
$$y_{n+1} = y_n + d_n\, v_n\, x_n \quad \to \sin(\theta) \tag{33.2-1h}$$
$$z_{n+1} = z_n - d_n\, A_n \quad \to 0 \tag{33.2-1i}$$

The scaling constant K is

$$K = \prod_{k=0}^{\infty} \sqrt{1 + 2^{-2k}} \tag{33.2-2a}$$
$$K = 1.64676025812106564836605122228229843565237672570102740 9\ldots \tag{33.2-2b}$$
$$\frac{1}{K} = 0.60725293500888125616944675250492826311239085215008977 24\ldots \tag{33.2-2c}$$

We note that K can be computed more efficiently as $K = \sqrt{2\,F(1/4)}$ where $F(z)$ is defined as

$$F(z) = \prod_{k=1}^{\infty} 1 + z^k \tag{33.2-3}$$

We use relation 14.4-6 on page 334, and relation 14.4-3b (pentagonal number theorem): $F(z) = P(z^2)/P(z)$ where

$$P(z) = 1 + \sum_{n=1}^{\infty} (-1)^n \left(z^{n(3n-1)/2} + z^{n(3n+1)/2}\right) \tag{33.2-4}$$

Using n terms of the sum gives a precision of about $3\,(n-1)^2$ bits:

```
? pent(z, n)= 1+sum(k=1,n, (-1)^k*(z^(k*(3*k-1)/2) + z^(k*(3*k+1)/2)));
? n=30; u=0.25; K=sqrt( 2 * pent(u^2,n)/pent(u,n) )
  1.64676025812106564836605122228229843565237672570102740 9
```

The CORDIC algorithm converges if $-r \leq z_0 \leq r$ where

$$r = \sum_{k=0}^{\infty} \arctan(2^{-k}) \tag{33.2-5a}$$

$$r = 1.74328662047234000350433765613641628581383118542820652 3\ldots \tag{33.2-5b}$$

$$r > \frac{\pi}{2} = 1.57079632\ldots \tag{33.2-5c}$$

With arguments x_0, y_0, z_0 one has

$$x \rightarrow K\left(x_0\,\cos(z_0) - y_0\,\sin(z_0)\right) \tag{33.2-6a}$$

$$y \rightarrow K\left(y_0\,\cos(z_0) + x_0\,\sin(z_0)\right) \tag{33.2-6b}$$

$$z \rightarrow 0 \tag{33.2-6c}$$

which, for $x_0 = 1/K, y_0 = 0, z_0 = \theta$ specializes to the computation as above.

A nice feature of the algorithm is that it also works backwards: initialize as above and use the same iteration with the slight modification that $d_n := -\operatorname{sign}(y_n)$, then

$$x \rightarrow K\sqrt{x_0^2 + y_0^2} \tag{33.2-7a}$$

$$y \rightarrow 0 \tag{33.2-7b}$$

$$z \rightarrow z_0 - \arctan\left(\frac{y_0}{x_0}\right) \tag{33.2-7c}$$

The algorithm can be derived by writing

$$\begin{bmatrix} x_{n+1} \\ y_{n+1} \end{bmatrix} = \begin{bmatrix} +\cos(d_n\,A_n) & -\sin(d_n\,A_n) \\ +\sin(d_n\,A_n) & +\cos(d_n\,A_n) \end{bmatrix} \begin{bmatrix} x_n \\ y_n \end{bmatrix} \tag{33.2-8}$$

and noting that (using $d_n = \pm 1$, so $\cos(d_n\,A_n) = \cos(A_n)$ and $\sin(d_n\,A_n) = d_n\,\sin(A_n)$)

$$\begin{bmatrix} x_{n+1} \\ y_{n+1} \end{bmatrix} = \cos(A_n) \begin{bmatrix} +1 & -d_n\,v_n \\ +d_n\,v_n & +1 \end{bmatrix} \begin{bmatrix} x_n \\ y_n \end{bmatrix} \tag{33.2-9}$$

where $v_n = 2^{-n}$. The CORDIC algorithm postpones the multiplications by $\cos(A_n)$. One has

$$\cos(A_n) = \cos\left(\arctan(2^{-n})\right) = \frac{1}{\sqrt{1 + 2^{-2n}}} \tag{33.2-10}$$

Thereby

$$K = 1/\prod_{k=0}^{\infty} \cos(A_n) = \prod_{k=0}^{\infty} \sqrt{1 + 2^{-2n}} \tag{33.2-11}$$

33.2.2 The linear case: multiplication and division

A slight variation gives a base-2 multiply-add algorithm:

$$A_n = 2^{-n} \tag{33.2-12a}$$

$$v_n = 2^{-n} \tag{33.2-12b}$$

$$d_n = \operatorname{sign}(z_n) \tag{33.2-12c}$$

$$x_{n+1} = x_n \tag{33.2-12d}$$

$$y_{n+1} = y_n + d_n\,v_n\,x_n \tag{33.2-12e}$$

$$z_{n+1} = z_n - d_n\,A_n \tag{33.2-12f}$$

then

$$x \;\to\; x_0 \tag{33.2-13a}$$
$$y \;\to\; y_0 + x_0\, z_0 \tag{33.2-13b}$$
$$z \;\to\; 0 \tag{33.2-13c}$$

Going backwards (replace relation 33.2-12c on the preceding page by $d_n := -\operatorname{sign}(y_n)$) gives an algorithm for division:

$$x \;\to\; x_0 \tag{33.2-14a}$$
$$y \;\to\; 0 \tag{33.2-14b}$$
$$z \;\to\; z_0 - \frac{y_0}{x_0} \tag{33.2-14c}$$

33.2.3 The hyperbolic case: *sinh* and *cosh*

n :	x_n	y_n	z_n	A_n
init	1.20749706	0.00000000	+1.00000000	+0.00000000
1:	1.20749706	0.60374853	+0.45069385	-0.54930614
2:	1.35843420	0.90562280	+0.19528104	-0.25541281
3:	1.47163705	1.07542707	+0.06962382	-0.12565721
4:	1.53885124	1.16740439	+0.00704225	-0.06258157
+4:	1.61181401	1.26358259	-0.05553931	-0.06258157
5:	1.57232706	1.21321340	-0.02427913	+0.03126017
6:	1.55337060	1.18864579	-0.00865286	+0.01562627
7:	1.54408430	1.17651008	-0.00084020	+0.00781265
8:	1.53948856	1.17047850	+0.00306606	+0.00390626
9:	1.54177465	1.17348532	+0.00111293	-0.00195312
10:	1.54292063	1.17499096	+0.00013637	-0.00097656
11:	1.54349436	1.17574434	-0.00035190	-0.00048828
12:	1.54320731	1.17536751	-0.00010776	+0.00024414
13:	1.54306383	1.17517913	+0.00001430	+0.00012207
+13:	1.54320729	1.17536749	-0.00010776	-0.00012207
14:	1.54313555	1.17527330	-0.00004673	+0.00006103
15:	1.54309968	1.17522621	-0.00001621	+0.00003051
∞:	1.54308063	1.17520119	+0.00000000	+0.00000000
	$= \cosh(1)$	$= \sinh(1)$	$= 0$	$= 0$

Figure 33.2-B: Numerical values occurring in the CORDIC computation of cosh(1) and sinh(1). Note that steps 4 and 13 are executed twice.

The versions presented so far can be unified as

$$v_n \;=\; 2^{-n} \tag{33.2-15a}$$
$$x_{n+1} \;=\; x_n - m\, d_n\, v_n\, y_n \tag{33.2-15b}$$
$$y_{n+1} \;=\; y_n + d_n\, v_n\, x_n \tag{33.2-15c}$$
$$z_{n+1} \;=\; z_n - d_n\, A_n \tag{33.2-15d}$$

where the linear case corresponds to $m = 0$ and $A_n = 2^{-n}$, the circular case to $m = 1$ and $A_n = \arctan(2^{-n})$. The forward direction ('rotation mode') is obtained by setting $d_n = \operatorname{sign}(z_n)$, the backward direction ('vectoring mode') by setting $d_n = -\operatorname{sign}(y_n)$.

Setting $m = -1$ gives a CORDIC algorithm that computes the hyperbolic sine and cosine or their inverses. The lookup table has to contain the values $\operatorname{arctanh}(2^{-n})$ for $n = 1, 2, 3, \ldots$, stored in the array `cordic_htab[]`. The algorithm needs a modification in order to converge: the iteration starts with index

one and some steps have to be executed twice. The sequence of the indices that need to be processed twice is $4, 13, 40, 121, ...$ ($i_0 = 4$, $i_{n+1} = 3\,i_n + 1$).

A sample implementation is given in [FXT: arith/cordic-hyp-demo.cc]:

```
1   void
2   cordic_hyp(double theta, double &s, double &c, ulong n)
3   {
4       double x = cordic_1Kp;
5       double y = 0;
6       double z = theta;
7       double v = 1.0;
8       // [PRINT]
9       ulong i = 4;
10      for (ulong k=1; k<n; ++k)
11      {
12          v *= 0.5;
13      again:
14          double d = ( z>=0 ? +1 : -1 );
15          double tx = x + d * v * y;
16          double ty = y + d * v * x;
17          double tz = z - d * cordic_htab[k];
18          x = tx;  y = ty;  z = tz;
19          // [PRINT]
20          if ( k==i )  { i=3*i+1; goto again; }
21      }
22      c = x;
23      s = y;
24  }
```

The values for first steps of the computation for the argument $\theta = z_1 = 1.0$ are given in figure 33.2-B.

The scaling constant corresponding to K is K', one has

$$K' = \prod_{k=1}^{\infty} \sqrt{1 - 2^{-2k}} \cdot \prod_{k=0}^{\infty} \sqrt{1 - 2^{-2i_k}} \tag{33.2-16a}$$

$$K' = 0.828159360960215627076198327759175146869453837690842529 1\ldots \tag{33.2-16b}$$

$$\frac{1}{K'} = 1.207497067763072128877721011310915836812783221769813422\ldots \tag{33.2-16c}$$

The duplicated indices appear twice in the product. The algorithm can be used for the computation of the exponential function using $\exp(x) = \sinh(x) + \cosh(x)$. The algorithm converges if $-r' \le z_1 \le r'$ where

$$r' = \sum_{k=1}^{\infty} \operatorname{arctanh}(2^{-k}) + \sum_{k=0}^{\infty} \operatorname{arctanh}(2^{-i_k}) \tag{33.2-17a}$$

$$r' = 1.118173015526503803610627556783092451806572942929536106\ldots \tag{33.2-17b}$$

With arguments x_1, y_1, z_1 one has

$$x \to K' \left(x_1 \cosh(z_1) + y_1 \sinh(z_1)\right) \tag{33.2-18a}$$

$$y \to K' \left(y_1 \cosh(z_1) + x_1 \sinh(z_1)\right) \tag{33.2-18b}$$

$$z \to 0 \tag{33.2-18c}$$

which, for $x_1 = 1/K', y_1 = 0, z_1 = \theta$ specializes to the computation as above.

The backward version ($d_n := -\operatorname{sign}(y_n)$) computes

$$x \to K' \sqrt{x_1^2 - y_1^2} \tag{33.2-19a}$$

$$y \to 0 \tag{33.2-19b}$$

$$z \to z_1 - \operatorname{arctanh}\left(\frac{y_1}{x_1}\right) \tag{33.2-19c}$$

For the computation of the natural logarithm use $\log(w) = 2 \operatorname{arctanh} \frac{w-1}{w+1}$. That is, start with $x_1 = w+1$ and $y_1 = w-1$, then $z \to \frac{1}{2} \log(w)$.

The computation of the square root $\sqrt{w}$ can be obtained by starting with $x_1 = w+1/4$ and $y_1 = w-1/4$ then $z \to K' \sqrt{w}$.

For further information see [14], [144], and chapter 6 of [205]. An algorithm working with complex numbers is given in [28].

Chapter 34

Recurrences and Chebyshev polynomials

This chapter presents algorithms and material concerning recurrences.

Firstly, several algorithms for recurrences, mostly for the case of constant coefficients, are given. Secondly, the Chebyshev polynomials are described. These are an important special case of a recurrence.

34.1 Recurrences

A sequence $[a_0, a_1, a_2, \dots]$ so that a *recurrence relation*

$$a_n \quad = \quad \sum_{j=1}^{k} m_j \, a_{n-j} \tag{34.1-1}$$

with given m_j holds for all a_j is called a k-th order *recurrence*. The recurrence is linear, homogeneous, with constant coefficients. The sequence is defined by both the recurrence relation and the first k elements.

An example, the second order recurrence relation $a_n = 1\,a_{n-1} + 1\,a_{n-2}$ together with $a_0 = 0$ and $a_1 = 1$ gives the Fibonacci numbers F_n, starting with $a_0 = 2$ and $a_1 = 1$ gives the Lucas numbers L_n:

```
  n : 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14
F(n): 0  1  1  2  3  5  8 13 21 34 55 89 144 233 377
L(n): 2  1  3  4  7 11 18 29 47 76 123 199 322 521 843
```

The *characteristic polynomial* of the recurrence relation 34.1-1 is given by

$$p(x) \quad = \quad x^k - \sum_{j=1}^{k} m_j \, x^{k-j} \tag{34.1-2}$$

The definition can be motivated by writing down the recurrence relation for the element with index $n = k$:

$$0 \quad = \quad a_k - \left(\sum_{j=1}^{k} m_j \, a_{k-j} \right) \tag{34.1-3}$$

34.1.1 Fast computation using matrix powers

For the recurrence defined by the recurrence relation

$$a_n \quad := \quad m_1 \, a_{n-1} + m_2 \, a_{n-2} \tag{34.1-4}$$

and the start a_0, a_1 use the relation

$$[a_0,\ a_1]\begin{bmatrix}0 & m_2\\ 1 & m_1\end{bmatrix}^k \quad = \quad [a_k,\ a_{k+1}] \tag{34.1-5}$$

for the fast computation of an individual element a_k. The algorithm is fast when powering algorithms (see section 27.6) are used.

Note that with two consecutive terms of the recurrence in the resulting vector it is easy two compute the following terms a_{k+1}, a_{k+2}, ... using the original recurrence relation.

The generalization is straightforward. For example, a recurrence $a_n = m_1\,a_{n-1} + m_2\,a_{n-2} + m_3\,a_{n-3}$ corresponds to

$$[a_0,\ a_1,\ a_2]\begin{bmatrix}0 & 0 & m_3\\ 1 & 0 & m_2\\ 0 & 1 & m_1\end{bmatrix}^k \quad = \quad [a_k,\ a_{k+1},\ a_{k+2}] \tag{34.1-6}$$

The matrix is the companion matrix of the characteristic polynomial $x^3 - (m_1\,x^2 + m_2\,x^1 + m_3\,x^0)$, see relation 40.5-1 on page 889. Note that the indexing of the m_k is different here.

Performance

The computations are fast. As an example we give the timing of the computation of a few sequence terms with large indices. The following calculations were carried out with exact arithmetic, the post-multiply with the float 1.0 renders the output readable:

```
? M=[0,1;1,1] \\ Fibonacci sequence
? #
   timer = 1 (on)
? ([0,1]*M^10000)[1]*1.0
  time = 1 ms.
  3.364476487643 E2089
? ([0,1]*M^100000)[1]*1.0
  time = 10 ms.
  2.597406934722 E20898
? ([0,1]*M^1000000)[1]*1.0
  time = 458 ms.
  1.953282128707 E208987
```

The powering algorithm can obviously be used also for polynomial recurrences such as for the Chebyshev polynomials $T_n(x)$:

```
? M=[0,-1;1,2*x]
  [0  -1]
  [1 2*x]
? for(n=0,5,print(n,": ",([1,x]*M^n)[1]))
  0: 1
  1: x
  2: 2*x^2 - 1
  3: 4*x^3 - 3*x
  4: 8*x^4 - 8*x^2 + 1
  5: 16*x^5 - 20*x^3 + 5*x
? p=([1,x]*M^1000)[1];
  time = 1,027 ms.
? poldegree(p)
  1000
? log(polcoeff(p,poldegree(p)))/log(10)
  300.728965668317  \\ The coefficient of x^1000 is a 301-digit number
```

With modular arithmetic the quantities remain bounded and the computations can be carried out for extreme large values of n. We use the modulus $m = 2^{1279} - 1$ and compute the $n = (m+1)/4$ element of the sequence $2, 4, 14, 52, \ldots$ where $a_n = 4\,a_{n-1} - a_{n-2}$:

```
? m=2^1279-1;  \\ a 1279-bit number
? log(m)/log(10)
  385.0173 \\ 306 decimal digits
? M=Mod([0,-1;1,4],m);  \\ all entries modulo m
?  component( ([2,4]*M^((m+1)/4))[1], 2)
  time = 118 ms.
  0
```

The result is zero which proves that m is prime, see section 37.11.4. Here is a one-liner that prints all exponents $e < 1000$ of Mersenne primes:

```
? forprime(e=3,1000,m=2^e-1;M=Mod([0,-1;1,4],m);if(0==(([2,4]*M^((m+1)/4))[1]),print1(" ", e)))
  3 5 7 13 17 19 31 61 89 107 127 521 607
```

The computation takes a few seconds only.

The connection of recurrences and matrix powers is investigated in [33].

34.1.2 Faster computation using polynomial arithmetic

The matrix power algorithm for computing the k-th element of a n-th order recursion involves proportional $\log_k$ multiplications of $n \times n$ matrices. As matrix multiplication (with the straightforward algorithm) is proportional n^3 the algorithm is not optimal for recursions of high order. Note that the matrix entries grow exponentially, so the asymptotics as given is valid only for computations with bounded values such as with modular arithmetic. We will see that the involved work can be brought down from $\log k \cdot n^3$ to $\log k \cdot n^2$ and even to $\log k \cdot n \cdot \log n$.

The characteristic polynomial for the recursion $a_n := 3\,a_{n-1} + 1\,a_{n-2} + 2\,a_{n-3}$ is

$$p(x) \quad = \quad x^3 - 3\,x^2 - 1\,x - 2 \tag{34.1-7}$$

We list the first few powers of the companion matrix M of $p(x)$:

$$M^0 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \qquad M^1 = \begin{bmatrix} 0 & 0 & 2 \\ 1 & 0 & 1 \\ 0 & 1 & 3 \end{bmatrix} \qquad M^2 = \begin{bmatrix} 0 & 2 & 6 \\ 0 & 1 & 5 \\ 1 & 3 & 10 \end{bmatrix} \qquad M^3 = \begin{bmatrix} 2 & 6 & 20 \\ 1 & 5 & 16 \\ 3 & 10 & 35 \end{bmatrix} \tag{34.1-8}$$

Note that each power is a left shifted version of its predecessor, only the rightmost colum is 'new'. Now compare the columns of the matrix powers to the first few values x^k modulo $p(x)$:

$$x^0 \bmod p(x) \quad = \quad 0\,x^2 + 0\,x + 1 \tag{34.1-9a}$$
$$x^1 \bmod p(x) \quad = \quad 0\,x^2 + 1\,x + 0 \tag{34.1-9b}$$
$$x^2 \bmod p(x) \quad = \quad 1\,x^2 + 0\,x + 0 \tag{34.1-9c}$$
$$x^3 \bmod p(x) \quad = \quad 3\,x^2 + 1\,x + 2 \tag{34.1-9d}$$
$$x^4 \bmod p(x) \quad = \quad 10\,x^2 + 5\,x + 6 \tag{34.1-9e}$$
$$x^5 \bmod p(x) \quad = \quad 35\,x^2 + 16\,x + 20 \tag{34.1-9f}$$

Observe that $x^k \bmod p(x)$ corresponds to the leftmost column of M^k.

We now turn the observation into an efficient algorithm. The main routines in this section take as arguments a vector v of initial values, a vector m of recursion coefficients and an index k. The vector $r = [a_k, a_{k+1}, \ldots, a_{k+n}]$ is returned. We compute the leftmost column of M^k as $z := x^k \bmod p(x)$ and compute a_k as the scalar product of z (as a vector) and v. Our main routine is:

```
1   frec(v, m, k)=
2   {
3       local(n, pc, pv, pp, px, r, t);
4       n = length(m);
5       if ( k<=n, return( recstep(v, m, k) ) );  \\ small indices by definition
6       pc = vec2charpol(m);
7       pp = Mod( x, pc );
8       px = pp^(k);
9       r = vector(n);
10      for (i=1, n,
11          t = component(px,2);
12          r[i] = sum(j=1,n, v[j]*polcoeff(t,j-1,x));
13          px *= pp;
14      );
15      return( r );
16  }
```

If only the value a_k is of interest, skip the computations in the final `for` loop for the values `i`> 1.

For small indices k the result is computed directly by definition, using the following auxiliary routine:

```
1   recstep(v, m, k)=
2   { /* update v by k steps according to the recursion coefficients in m */
3       local(n,r);
4       if ( k<=0, return(v) );  \\ negative k is forbidden
5       n = length(m);
6       r = vector(n);
7       for (i=1, k,
8           for (j=1, n-1, r[j]=v[j+1] ); \\ shift left
9           r[n] = sum(j=1,n, m[n+1-j]*v[j]);  \\ new element (convolution)
10          v = r;
11      );
12      return( r );
13  }
```

The auxiliary routine used to compute the characteristic polynomial corresponding to the vector m is:

```
1   vec2charpol(m)=
2   { /* return characteristic polynomial for the recursion coefficients in m */
3       local(d,p);
4       d = length(m);
5       p = x^d - Pol(m,x);
6       return( p );
7   }
```

The computation of the k-th element of a n-term recurrence involves proportional $\log k$ modular polynomial multiplications. Thereby the total cost is $\log k \cdot M(n)$ where $M(n)$ is the cost of the multiplication of two polynomials of degree n. That is, the cost is $\log k \cdot n^2$ when usual polynomial multiplication is used, and $\log k \cdot n \cdot \log n$ if an FFT scheme is applied.

The matrix power algorithm, restated for the argument structure defined above, can be implemented as:

```
1   mrec(v, m, k)=
2   {
3       local(p,M);
4       p = vec2charpol(m);
5       M = matcompanion(p);
6       M = M^k;
7       return ( v * M );
8   }
```

All main routines can be used with symbolic values:

```
? frec([a0,a1],[m1,m2],3)
  [m2*m1*a0 + (m1^2 + m2)*a1, (m2*m1^2 + m2^2)*a0 + (m1^3 + 2*m2*m1)*a1]
? mrec([a0,a1],[m1,m2],3)
  [m2*m1*a0 + (m1^2 + m2)*a1, (m2*m1^2 + m2^2)*a0 + (m1^3 + 2*m2*m1)*a1]
? recstep([a0,a1],[m1,m2],3)
  [m2*m1*a0 + (m1^2 + m2)*a1, (m2*m1^2 + m2^2)*a0 + (m1^3 + 2*m2*m1)*a1]
```

Performance

We check the performance (suppressing output):

```
? k=10^5;
? recstep([0,1],[1,1],k);
  time = 2,811 ms.  \\ time linear in k
? mrec([0,1],[1,1],k);
  time = 10 ms. \\ time linear in log(k)
? frec([0,1],[1,1],k);
  time = 4 ms. \\ time linear in log(k)
```

The relative performance of the routine `frec()` and `mrec()` differs more with higher orders n of the recurrence, we use $n = 10$:

```
? n=10; v=vector(n); v[n]=1; m=vector(n,j,1); k=10^5;  \\ tenth order recurrence
? mrec(v,m,k);
  time = 2,813 ms.
? f=frec(v,m,k);
  time =  159 ms.
```

```
? log(f)/log(10.0)
  [30078.67, 30078.97, 30079.27, 30079.58, 30079.88, 30080.18, \
  30080.48, 30080.78, 30081.08, 30081.38] \\ about 30k decimal digits each
```

Somewhat surprisingly, we see a performance gain greater than n even though the computations were done using integers. Finally, we repeat the computations modulo $p = 2^{521} - 1$ for $k = 10^{30}$:

```
?  n=10; v=vector(n); v[n]=1; m=vector(n,j,1); k=10^30;
?  p=2^521-1; v=Mod(v,p); m=Mod(m,p);
?  mrec(v,m,k);
   time = 312 ms.
?  frec(v,m,k);
   time = 14 ms.
```

That the performance gain with integers is not smaller than with modular arithmetic can be motivated by the fact that the quantities in both algorithms grow with the same rate. Now at each step the performance ratio should approximately equal n. Thereby the algorithms perform with the same ratio.

The computational advantage of powering modulo the characteristic polynomial versus matrix powering has been pointed out 1994 by Brent [66, p.392] (page 4 of the preprint).

34.1.3 Inhomogeneous recurrences

The fast algorithms for the computation of recurrences do only work with *homogeneous* recurrences as defined by relation 34.1-1 on page 657. A *inhomogeneous* recurrence is defined by a relation

$$a_n \quad = \quad \sum_{j=1}^{k} m_j \, a_{n-j} + P(n) \tag{34.1-10}$$

where $P(n)$ is a nonzero polynomial in n. We will show how to transform a inhomogeneous recurrence into a homogeneous recurrence of greater order.

34.1.3.1 Recurrence relations with a constant

In case a constant is to be added in an k-th order relation, one can use a recurrence of order $k+1$. From the recurrence relation $a_n = m_1 \, a_{n-1} + m_2 \, a_{n-2} + \ldots + m_k \, a_{n-k} + C$ subtract a shifted version $a_{n-1} = m_1 \, a_{n-2} + m_2 \, a_{n-3} + \ldots + m_k \, a_{n-k-1} + C$ to obtain $a_n = (m_1 + 1) \, a_{n-1} + (m_2 - m_1) \, a_{n-2} + \ldots + (m_k - m_{k-1}) \, a_{n-k}$.

An example should make the idea clear: with $a_n = 34 \, a_{n-1} - a_{n-2} + 2$ subtract a shifted version $a_{n-1} = 34 \, a_{n-2} - a_{n-3} + 2$ to obtain $a_n = 35 \, a_{n-1} - 35 \, a_{n-2} + a_{n-3}$. Setting $a_0 = 1$, $a_1 = 36$ we get, using the original relation

```
? n=7;
? ts=vector(n); ts[1]=1; ts[2]=36;
? for(k=3,n,ts[k]=34*ts[k-1]-ts[k-2]+2);
? ts
  [1, 36, 1225, 41616, 1413721, 48024900, 1631432881]
```

and, using the relation without constant,

```
? ts=vector(n); ts[1]=1; ts[2]=36; ts[3]=34*ts[2]-ts[1]+2;
? for(k=4,n,ts[k]=35*ts[k-1]-35*ts[k-2]+ts[k-3]);
? ts
  [1, 36, 1225, 41616, 1413721, 48024900, 1631432881]
```

34.1.3.2 The general case

If the recurrence is of the form

$$a_n \quad = \quad m_1 \, a_{n-1} + m_2 \, a_{n-2} + \ldots + m_k \, a_{n-k} + P(n) \tag{34.1-11}$$

where $P(n)$ is a polynomial of degree d in n then a homogeneous recurrence of order $k+d+1$

$$a_n \quad = \quad M_1\, a_{n-1} + M_2\, a_{n-2} + \ldots + M_{k+d+1}\, a_{n-k-d-1} \tag{34.1-12}$$

can be obtained by repeatedly subtracting a shifted relation.

The following pari/gp routine takes as input a vector of the multipliers m_i ($i = 1, \ldots, k$) and a polynomial of degree d in n. It returns a homogeneous recurrence relation as a vector $[M_1, \ldots, M_{k+d+1}]$:

```
1   ihom2hom(m, p)=
2   {
3       local(d, M, k);
4       if ( p==0, return(m) );
5       d = poldegree(p, 'n);
6       k = length(m);
7       M = vector(k+d+1);
8       for (j=1, k, M[j]=m[j]);
9       for (s=1, d+1,
10          M[1] += 1;  \\ left hand side
11          for (j=2, k+s, M[j] -= m[j-1]; );
12          m = M;
13      );
14      return(M);
15  }
```

In order to verify the output we use a (slow) routine that directly computes the values of an inhomogeneous recurrence:

```
1   ihom(v, m, k, p)=
2   {
3       local(n, r);
4       if ( k<=0, return(v[1]) );
5       n = length(m);
6       r = vector(n);
7       for (i=1, k,
8           for (j=1, n-1, r[j]=v[j+1] ); \\ shift left
9           r[n] = sum(j=1,n, m[n+1-j]*v[j]);  \\ new element (convolution)
10          r[n] += subst(p, 'n, i+n-1);  \\ add inhomogeneous term
11          v = r;
12      );
13      return( r[1] );
14  }
```

We use the recurrence relation $a_n = 3\,a_{n-1} + 2a_{n-2} + (n^3 - n^2 - 7)$. We compute the homogeneous equivalent (intermediate values of M added):

```
? m=[3,+2];p=n^3-n^2-7;
? M=ihom2hom(m,p)
    [3, 2, 0, 0, 0, 0]
    [4, -1, -2, 0, 0, 0]
    [5, -5, -1, 2, 0, 0]
    [6, -10, 4, 3, -2, 0]
    [7, -16, 14, -1, -5, 2]
  [7, -16, 14, -1, -5, 2]  \\ a_n = 7*a_{n-1} - 16*a_{n-2} + 14*a_{n-3} +- ...
```

We can compute the first few values for the sequence starting with $a_0 = 2$, $a_1 = 5$ by the direct method:

```
? v=[2,5];
? for(k=0,9,print(k,":  ",ihom(v,m,k,p)));
  0:  2
  1:  5
  2:  16
  3:  69
  4:  280
  5:  1071
  6:  3946
  7:  14267
  8:  51134
  9:  182577
```

A vector of start values and the homogeneous equivalent allow the fast computation using the powering algorithms:

```
? V=vector(length(M),j,ihom(v,m,j-1,p))
  [2, 5, 16, 69, 280, 1071]
? for(k=0,9,print(k,":  ",frec(V,M,k)[1]));
 [- same output as with direct computation -]
```

The computation of $a_{10,000}$ now takes less than a second:

```
? z=frec(V,M,10^5)[1];  \\ result computed in 156 ms.
? 1.0*z
  1.72279531330182 E55164
? z=ihom(v,m,10^5,p);  \\ result computed in 6,768 ms.
```

34.1.4 Recurrence relations for subsequences

34.1.4.1 Two term recurrences

The recurrence for the subsequence of every k-th element of a two term recurrence $a_n = \alpha\, a_{n-1} + \beta\, a_{n-2}$ can be obtained as follows. Write

$$\begin{aligned}
a_{n+0} &= A_0\, a_n + B_0\, a_{n-0} &&= 2\, a_n \;-\; 1\, a_{n-0} && \text{(34.1-13a)}\\
a_{n+1} &= A_1\, a_n + B_1\, a_{n-1} &&= \alpha\, a_n \;+\; \beta\, a_{n-1} && \text{(34.1-13b)}\\
a_{n+2} &= A_2\, a_n + B_2\, a_{n-2} &&= (\alpha^2 + 2\beta)\, a_n \;-\; \beta^2\, a_{n-2} && \text{(34.1-13c)}\\
a_{n+3} &= A_3\, a_n + B_3\, a_{n-3} &&= (\alpha^3 + 3\alpha\beta)\, a_n \;+\; \beta^3\, a_{n-3} && \text{(34.1-13d)}\\
a_{n+4} &= A_4\, a_n + B_4\, a_{n-4} &&= (\alpha^4 + 4\alpha^2\beta + 2\beta^2)\, a_n \;-\; \beta^4\, a_{n-4} && \text{(34.1-13e)}\\
a_{n+k} &= A_k\, a_n + B_k\, a_{n-k} && && \text{(34.1-13f)}
\end{aligned}$$

We have $a_n = A_k\, a_{n-k} + B_k\, a_{n-2k}$ where $A_0 = 2$, $A_1 = \alpha$ and $A_{k+1} = \alpha\, A_k + \beta\, A_{k-1}$ (and $B_k = -(-\beta)^k$). That is, the first coefficient A_k of the recursion relations for the subsequences can be computed by the original recurrence relation. For efficient computation use

$$[A_k,\ A_{k+1}] \;=\; [2,\ \alpha] \begin{bmatrix} 0 & \beta \\ 1 & \alpha \end{bmatrix}^k \qquad \text{(34.1-14)}$$

A closed form for A_k in terms of Chebyshev polynomials is given in [31, item 14]:

$$A_k \;=\; 2\,(-\beta)^{k/2}\, T_k\left(\alpha/\sqrt{-4\,\beta}\right) \qquad \text{(34.1-15)}$$

A simple example, let F_n and L_n denote the n-th Fibonacci and Lucas number, respectively. Then $\alpha = \beta = 1$ and

$$[A_k,\ A_{k+1}] \;=\; [2,\ 1] \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}^k \;=\; [L_k,\ L_{k+1}] \qquad \text{(34.1-16)}$$

That is

$$F_{k\,n+e} \;=\; L_k\, F_{k\,(n-1)+e} - (-1)^k\, F_{k\,(n-2)+e} \qquad \text{(34.1-17)}$$

where $k \in \mathbb{Z}$ and $e \in \mathbb{Z}$. The variable e expresses the shift invariance of the relation.

34.1.4.2 Recurrences of order n

For the stride-s recurrence relations of order n the following may be the most straightforward algorithm. Let $p(x)$ be the characteristic polynomial of the recurrence and M its companion matrix. Then the characteristic polynomial of M^s corresponds to the recurrence relation of the stride-s subsequence.

```
1  recsubseq(n, s, m=0)=
2  { /* Return vector coefficients of the stride-s subsequence
3     * of the n-th order linear recurrence.
4     */
5      local(p, M, z, r);
6      if ( 0==m,
```

```
 7          m = vector(n,j,eval(Str("m" j)));  \\ use symbols m_j
 8          , /* else */
 9          n = length(m);  \\ m given
10        );
11       p = vec2charpol(m);
12       M = matcompanion(p);
13       z = x^n-charpoly(M^s);
14       r = vector(n,j,polcoeff(z,n-j,x));
15       return( r );
16   }
```

For the second order recurrence we get what we have already seen for $s = 0, \ldots, 4$:

```
? m=[a,b];
? for(s=-2,5,print(s,":  ",recsubseq(0,s,m)););
  -2:  [1/b^2*a^2 + 2/b,   -1/b^2]
  -1:  [-1/b*a,   1/b]
   0:  [2,   -1]
   1:  [a,   b]
   2:  [a^2 + 2*b,   -b^2]
   3:  [a^3 + 3*b*a,   b^3]
   4:  [a^4 + 4*b*a^2 + 2*b^2,   -b^4]
   5:  [a^5 + 5*b*a^3 + 5*b^2*a,   b^5]
```

For the third order recurrence we get:

```
? m=[a,b,c];
? for(s=-2,5,print(s,":  ",recsubseq(0,s,m)););
  -2:  [2/-c*a - 1/-c^2*b^2,   1/-c^2*a^2 + 2/-c^2*b,   1/c^2]
  -1:  [-1/c*b,   1/-c*a,   1/c]
   0:  [3,   -3,   1]
   1:  [a,   b,   c]
   2:  [a^2 + 2*b,   2*c*a - b^2,   c^2]
   3:  [a^3 + 3*b*a + 3*c,   -3*c*b*a + (b^3 - 3*c^2),   c^3]
   4:  [a^4 + 4*b*a^2 + 4*c*a + 2*b^2, \
        -2*c^2*a^2 + 4*c*b^2*a + (-b^4 + 4*c^2*b), \
        c^4]
   5:  [a^5 + 5*b*a^3 + 5*c*a^2 + 5*b^2*a + 5*c*b, \
        5*c^2*b*a^2 + (-5*c*b^3 + 5*c^3)*a + (b^5 - 5*c^2*b^2), \
        c^5]
```

34.1.5 Generating functions for recurrences

A generating function for a recurrence has a power series where the k-th coefficient equals the k-th term of the recurrence. For example, for the Fibonacci- and Lucas numbers:

$$\frac{x}{1-x-x^2} = 0 + x + x^2 + 2x^3 + 3x^4 + 5x^5 + 8x^6 + 13x^7 + 21x^8 + 34x^9 + \ldots \quad (34.1\text{-}18a)$$

$$= \sum_{k=0}^{\infty} F_k\, x^k \quad (34.1\text{-}18b)$$

$$\frac{2-x}{1-x-x^2} = 2 + x + 3x^2 + 4x^3 + 7x^4 + 11x^5 + 18x^6 + 29x^7 + 47x^8 + \ldots \quad (34.1\text{-}18c)$$

$$= \sum_{k=0}^{\infty} L_k\, x^k \quad (34.1\text{-}18d)$$

In general, for a recurrence $a_n = \sum_{k=1}^{K} m_k\, a_{n-k}$ with given a_0, a_1, …, a_K one has

$$\frac{\sum_{j=0}^{K-1} b_j\, x^j}{1 - \sum_{j=1}^{K} m_j\, x^j} = \sum_{j=0}^{\infty} a_j\, x^j \quad (34.1\text{-}19a)$$

where the denominator is the reciprocal polynomial of the characteristic polynomial and

$$\begin{aligned} b_0 &= a_0 && \text{(34.1-20a)}\\ b_1 &= a_1 - (a_0\, m_1) && \text{(34.1-20b)}\\ b_2 &= a_2 - (a_0\, m_2 + a_1\, m_1) && \text{(34.1-20c)}\\ b_3 &= a_3 - (a_0\, m_3 + a_1\, m_2 + a_2\, m_1) && \text{(34.1-20d)}\\ b_k &= a_k - \sum_{j=0}^{k-1} a_j\, m_{k-j} && \text{(34.1-20e)} \end{aligned}$$

As an example we choose the sequence

$$[0, 0, 1, 1, 2, 4, 7, 13, 24, 44, 81, 149, 274, \ldots] \qquad \text{(34.1-21)}$$

with the recurrence relation $a_n = a_{n-1} + a_{n-2} + a_{n-2}$:

```
? a=[0,0,1]~;
? m=[1,1,1]~;
? K=length(m);
? b=vector(K, k, a[k]-sum(j=0,k-2, a[j+1]*m[k-j-1]))
  [0, 0, 1]
? pb=sum(j=0,K-1,b[j+1]*x^j)
  x^2
? pr=1-sum(k=1,K,m[k]*x^k) \\ reciprocal of charpoly
  -x^3 - x^2 - x + 1
? gen=pb/pr  \\ the generating function
  x^2/(-x^3 - x^2 - x + 1)
? t=taylor(gen, x)
  x^2 + x^3 + 2*x^4 + 4*x^5 + 7*x^6 + 13*x^7 + 24*x^8 + 44*x^9 + 81*x^10
  + 149*x^11 + 274*x^12 + 504*x^13 + 927*x^14 + 1705*x^15 + 3136*x^16 + O(x^17)
? t=truncate(t);
? for(j=0,poldegree(t),print1("  ",polcoeff(t,j)))
  0  0  1  1  2  4  7  13  24  44  81  149  274  504  927  1705  3136
```

Note that the denominator is the reciprocal of the characteristic polynomial. The general form of the expressions for a two term linear recurrence can be obtained using symbols:

```
? a=[a0,a1]~;
? m=[m1,m2]~;
? K=length(m);
? b=vector(K,k,a[k]-sum(j=0,k-2,a[j+1]*m[k-j-1]))
  [a0, -m1*a0 + a1]
? pb=sum(j=0,K-1,b[j+1]*x^j)
  (-m1*a0 + a1)*x + a0
? pr=1-sum(k=1,K,m[k]*x^k)
  -m2*x^2 - m1*x + 1
? gen=pb/pr  \\ the generating function
  ((-m1*a0 + a1)*x + a0)/(-m2*x^2 - m1*x + 1)
? t=taylor(gen,x);
? t=truncate(t);
? for(j=0,poldegree(t),print(j,":  ",polcoeff(t,j)))
  0:  a0
  1:  a1
  2:  m2*a0 + m1*a1
  3:  m2*m1*a0 + (m1^2 + m2)*a1
  4:  (m2*m1^2 + m2^2)*a0 + (m1^3 + 2*m2*m1)*a1
  5:  (m2*m1^3 + 2*m2^2*m1)*a0 + (m1^4 + 3*m2*m1^2 + m2^2)*a1
  6:  (m2*m1^4 + 3*m2^2*m1^2 + m2^3)*a0 + (m1^5 + 4*m2*m1^3 + 3*m2^2*m1)*a1
  7:  (m2*m1^5 + 4*m2^2*m1^3 + 3*m2^3*m1)*a0 + (m1^6 + 5*m2*m1^4 + 6*m2^2*m1^2 + m2^3)*a1
```

34.1.6 Binet forms for recurrences

A closed form expression for the Fibonacci numbers is

$$F_n = \frac{1}{\sqrt{5}} \left[\left(\frac{1+\sqrt{5}}{2} \right)^n - \left(\frac{1-\sqrt{5}}{2} \right)^n \right] \qquad \text{(34.1-22)}$$

For a two-term recurrence $a_n = m_1\, a_{n-1} + m_2\, a_{n-2}$ a closed form solution is given by

$$a_n \quad = \quad \frac{1}{w}\,[(a_1 - a_0\, M)\, P^n - (a_1 - a_0\, P)\, M^n] \tag{34.1-23a}$$

where $w = \sqrt{m_1^2 + 4\, m_2}$, $P = (m_1 + w)/2$ and $M = (m_1 - w)/2$.

In general such formulas can be obtained as exemplified using a three-term recurrence: let $a_n = m_1\, a_{n-1} + m_2\, a_{n-2} + m_3\, a_{n-3}$, its characteristic polynomial is $p(x) = x^3 - (m_1\, x^2 + m_2\, x + m_3)$. Let r_0, r_1, r_2 be the roots of $p(x)$, then $a_n = c_0\, r_0^n + c_1\, r_1^n + c_2\, r_2^n$ if c_0, c_1, c_2 satisfy

$$a_0 \quad = \quad c_0 + \quad c_1 + \quad c_2 \tag{34.1-24a}$$

$$a_1 \quad = \quad r_0\, c_0 + r_1\, c_1 + r_2\, c_2 \tag{34.1-24b}$$

$$a_2 \quad = \quad r_0^2\, c_0 + r_1^2\, c_1 + r_2^2\, c_2 \tag{34.1-24c}$$

That is, we have to solve the matrix equation $Z \cdot c = a$ for the vector c where a is the vector of starting values and

$$Z \quad = \quad \begin{bmatrix} 1 & 1 & 1 \\ r_0 & r_1 & r_2 \\ r_0^2 & r_1^2 & r_2^2 \end{bmatrix} \tag{34.1-25}$$

Verification with the three term recurrence $a_n = a_{n-1} + a_{n-2} + a_{n-2}$ starting with $a_0 = a_1 = 0$ and $a_2 = 1$:

```
? a=[0,0,1]~;
? m=[1,1,1]~;
? K=length(m);
? p=x^K-sum(k=1,K,m[k]*x^(K-k)) \\ characteristic polynomial
  x^3 - x^2 - x - 1
? r=(polroots(p))
  [1.8392867,   -0.419643 - 0.606290*I,  -0.419643 + 0.6062907*I]~
? Z=matrix(K,K,ri,ci,r[ci]^(ri-1))
  [1 1 1]
  [1.839286   -0.4196433 - 0.6062907*I   -0.4196433 + 0.6062907*I]
  [3.382975   -0.1914878 + 0.5088517*I   -0.1914878 - 0.5088517*I]

? c=matsolve(Z,a)
  [0.1828035 + 1.8947 E-20*I,  -0.09140176 - 0.3405465*I,  -0.0914017 + 0.3405465*I]~
? norm(Z*c-a) \\ check solution
  [1.147 E-39,  6.795 E-39,  3.673 E-39]~
? seq(n)=sum(k=0,K-1,c[k+1]*r[k+1]^n)
? for(n=0,20,print1(" ",round(seq(n))))
  0 0 1 1 2 4 7 13 24 44 81 149 274 504 927 1705 3136 5768 10609 19513 35890
```

The method fails if the characteristic polynomial has multiple roots because then the matrix Z is singular.

34.1.6.1 The special case $c_k = 1$

Let $p(x)$ be the characteristic polynomial of a recurrence, with roots r_i: $p(x) = \prod_k (x - r_k)$. We want to determine the generating function for the recurrence such that $a_j = \sum_k r_k^j$ (that is, all constants c_k are one). For the reciprocal polynomial h of p we have $h(x) = \prod_k (1 - r_k\, x_k)$, and (using the product rule for differentiation)

$$h'(x) \quad = \quad h(x) \sum_k \frac{-r_k}{1 - r_k\, x} \tag{34.1-26}$$

With $r/(1 - r\, x) = \sum_{j \geq 0} r^{j+1}\, x^j$ we find that

$$-\frac{h'(x)}{h(x)} \quad = \quad \sum_{j \geq 0} \left(\sum_k r_k^{j+1} \right) x^j \tag{34.1-27}$$

That is $a_j = \sum_k r_k^{j+1}$, and $c_k = 1$ for all k. The relation is the key to the fast computation of the trace vector in finite fields, see relation 40.3-6 on page 885.

34.1.6.2 Binet form with multiple roots of the characteristic polynomial

When the characteristic polynomial has multiple roots the Binet form has coefficients that are polynomials in n. For example, for the characteristic polynomials $p(x) = (x - r_0)^3 (x - r_1)$ the Binet from would be $a_n = (c_0 + n\,d_0 + n^2\,e_0)\,r_0^n + c_1\,r_1^n$. With $n = 0$, 1, and 2 we obtain the system of equations

$$a_0 = (c_0 + 0\,d_0 + 0^2\,e_0) + c_1 \tag{34.1-28a}$$
$$a_1 = r_0\,(c_0 + 1\,d_0 + 1^2\,e_0) + r_1\,c_1 \tag{34.1-28b}$$
$$a_2 = r_0^2\,(c_0 + 2\,d_0 + 2^2\,e_0) + r_1^2\,c_1 \tag{34.1-28c}$$

In general, the coefficient of the power of the k-th root r_k in the Binet form must be a polynomial of degree $m_k - 1$ where m_k is the multiplicity of r_k.

34.1.7 Logarithms of generating functions *

A seemingly mysterious relation for the generating function of the Fibonacci numbers

$$f(x) := \frac{1}{1 - x - x^2} = 1 + x + 2x^2 + 3x^3 + 5x^4 + 8x^5 + \ldots = \sum_{k=0}^{\infty} F_{k+1}\,x^k \tag{34.1-29}$$

is

$$\log(f(x)) = x + \frac{1}{2}\,3\,x^2 + \frac{1}{3}\,4\,x^3 + \frac{1}{4}\,7\,x^4 + \frac{1}{5}\,11\,x^5 + \ldots = \sum_{k=1}^{\infty} \frac{1}{k}\,L_k\,x^k \tag{34.1-30}$$

where L_k are the Lucas numbers. Similarly,

$$g(x) := \frac{1}{1 - 2\,x - x^2} = 1 + 2x + 5x^2 + 12x^3 + 29x^4 + 70x^5 + 169x^6 + \ldots \tag{34.1-31a}$$
$$\log(g(x)) = 2\left[x + \frac{1}{2}\,3\,x^2 + \frac{1}{3}\,7\,x^3 + \frac{1}{4}\,17\,x^4 + \frac{1}{5}\,41\,x^5 + \frac{1}{6}\,99\,x^6 + \ldots\right] \tag{34.1-31b}$$

Now set $f(x) =: \frac{1}{h(x)}$, then

$$\frac{d}{d\,x}\log(f(x)) = \frac{d}{d\,x}\log\left(\frac{1}{h(x)}\right) = -\frac{h'(x)}{h(x)} \tag{34.1-32}$$

The expression $\frac{h'(x)}{h(x)}$ is again the generating function of a recurrence and formal integration of the Taylor series terms gives the factors $\frac{1}{k}$. The observation is a special case of the algorithm for the computation of the logarithm for powers series given in section 31.3 on page 626.

34.2 Chebyshev polynomials

The *Chebyshev polynomials* of the first (T) and second (U) kind can be defined by the functions

$$T_n(x) = \cos\left[n\,\arccos(x)\right] \tag{34.2-1a}$$
$$U_n(x) = \frac{\sin\left[(n+1)\,\arccos(x)\right]}{\sqrt{1 - x^2}} \tag{34.2-1b}$$

For integral n both of them are polynomials. The first few polynomials are given in figure 34.2-A (first kind) and figure 34.2-B (second kind).

$$
\begin{aligned}
T_{-n}(x) &= T_n(x) \\
T_{-1}(x) &= x \\
T_0(x) &= 1 \\
T_1(x) &= x \\
T_2(x) &= 2\,x^2 - 1 \\
T_3(x) &= 4\,x^3 - 3\,x \\
T_4(x) &= 8\,x^4 - 8\,x^2 + 1 \\
T_5(x) &= 16\,x^5 - 20\,x^3 + 5\,x \\
T_6(x) &= 32\,x^6 - 48\,x^4 + 18\,x^2 - 1 \\
T_7(x) &= 64\,x^7 - 112\,x^5 + 56\,x^3 - 7\,x \\
T_8(x) &= 128\,x^8 - 256\,x^6 + 160\,x^4 - 32\,x^2 + 1 \\
T_9(x) &= 256\,x^9 - 576\,x^7 + 432\,x^5 - 120\,x^3 + 9\,x \\
T_{10}(x) &= 512\,x^{10} - 1280\,x^8 + 1120\,x^6 - 400\,x^4 + 50\,x^2 - 1 \\
T_{11}(x) &= 1024\,x^{11} - 2816\,x^9 + 2816\,x^7 - 1232\,x^5 + 220\,x^3 - 11\,x
\end{aligned}
$$

Figure 34.2-A: The first few Chebyshev polynomials of the first kind.

$$
\begin{aligned}
U_{-n}(x) &= -U_{n-2}(x) \\
U_{-2}(x) &= -1 \\
U_{-1}(x) &= 0 \\
U_0(x) &= 1 \\
U_1(x) &= 2\,x \\
U_2(x) &= 4\,x^2 - 1 \\
U_3(x) &= 8\,x^3 - 4\,x \\
U_4(x) &= 16\,x^4 - 12\,x^2 + 1 \\
U_5(x) &= 32\,x^5 - 32\,x^3 + 6\,x \\
U_6(x) &= 64\,x^6 - 80\,x^4 + 24\,x^2 - 1 \\
U_7(x) &= 128\,x^7 - 192\,x^5 + 80\,x^3 - 8\,x \\
U_8(x) &= 256\,x^8 - 448\,x^6 + 240\,x^4 - 40\,x^2 + 1 \\
U_9(x) &= 512\,x^9 - 1024\,x^7 + 672\,x^5 - 160\,x^3 + 10\,x \\
U_{10}(x) &= 1024\,x^{10} - 2304\,x^8 + 1792\,x^6 - 560\,x^4 + 60\,x^2 - 1 \\
U_{11}(x) &= 2048\,x^{11} - 5120\,x^9 + 4608\,x^7 - 1792\,x^5 + 280\,x^3 - 12\,x
\end{aligned}
$$

Figure 34.2-B: The first few Chebyshev polynomials of the second kind.

One has

$$T_n(x) = \frac{n}{2} \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k \frac{(n-k-1)!}{k!\,(n-2k)!} (2x)^{n-2k} \tag{34.2-4a}$$

$$= \frac{n}{2} \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k \frac{1}{n-k} \binom{n-k}{k} (2x)^{n-2k} \tag{34.2-4b}$$

$$= \sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2\,k} x^{n-2k} (x^2-1)^k \tag{34.2-4c}$$

and

$$U_n(x) = \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k \frac{(n-k)!}{k!\,(n-2k)!} (2x)^{n-2k} \tag{34.2-5a}$$

$$= \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k \binom{n-k}{k} (2x)^{n-2k} \tag{34.2-5b}$$

$$= \sum_{k=0}^{\lfloor n/2+1 \rfloor} \binom{n+1}{2\,k+1} x^{n-2k} (x^2-1)^k \tag{34.2-5c}$$

The indexing of U seems to be slightly unfortunate, having $U_0 = 0$ would render many of the relations for the Chebyshev polynomials more symmetric.

The $n+1$ extrema of $T_n(x)$ are located at the points $x_k = \cos \frac{k\,\pi}{n}$ where $k = 0, 1, 2, \ldots, n$ and $-1 \leq x_k \leq +1$, which can be seen from the definition. The values at those points are ± 1. The n zeros lie at $x_k = \cos \frac{(k-1/2)\,\pi}{n}$ where $k = 1, 2, 3, \ldots, n$.

The expansion of x^n in terms of Chebyshev polynomials of the first kind is, for n even,

$$x^n = \frac{1}{2^n} \binom{n}{n/2} + \frac{1}{2^{n-1}} \sum_{k=0}^{n/2-1} \binom{n}{k} T_{n-2k}(x) \tag{34.2-6a}$$

and, for odd n,

$$x^n = \frac{1}{2^{n-1}} \sum_{k=0}^{(n-1)/2} \binom{n}{k} T_{n-2k}(x) \tag{34.2-6b}$$

For the Chebyshev polynomials of the first kind one has

$$T_n\left(\frac{x+1/x}{2}\right) = \frac{x^n + 1/x^n}{2} \tag{34.2-7}$$

This relation can be used to find a solution of $T_n(x) = z$ directly. Indeed

$$x = \frac{R_n + 1/R_n}{2} \quad \text{where} \quad R_n := \left(z + \sqrt{z^2-1}\right)^{1/n} \tag{34.2-8}$$

is a solution which can be chosen to be real if $z \in \mathbb{R}$ and $z > 1$. Thereby we have the closed form expression

$$T_n(z) = \frac{r^n + r^{-n}}{2} \quad \text{where} \quad r := \left(z + \sqrt{z^2-1}\right) \tag{34.2-9}$$

34.2.1 Recurrence relation, generating functions, and the composition law

Both types of Chebyshev polynomials obey the same recurrence (omitting the argument x)

$$N_n = 2\,x\,N_{n-1} - N_{n-2} \tag{34.2-10}$$

where N can be either symbol, T or U. Recurrence relations for subsequences are:

$$\begin{aligned}
N_{n+1} &= [2\,x]\cdot N_n - N_{n-1} && \text{(34.2-11a)}\\
N_{n+2} &= \left[2\,(2x^2-1)\right]\cdot N_n - N_{n-2} && \text{(34.2-11b)}\\
N_{n+3} &= \left[2\,(4x^3-3x)\right]\cdot N_n - N_{n-3} && \text{(34.2-11c)}\\
N_{n+4} &= \left[2\,(8x^4-8x^2+1\right]\cdot N_n - N_{n-4} && \text{(34.2-11d)}\\
N_{n+5} &= \left[2\,(16x^5-20x^3+5x)\right]\cdot N_n - N_{n-5} && \text{(34.2-11e)}\\
N_{n+s} &= [2\,T_s(x)]\cdot N_n - N_{n-s} && \text{(34.2-11f)}
\end{aligned}$$

The generating functions are

$$\begin{aligned}
\frac{1-x\,t}{1-2\,x\,t+t^2} &= \sum_{n=0}^{\infty} t^n\,T_n(x) && \text{(34.2-12a)}\\
\frac{1}{1-2\,x\,t+t^2} &= \sum_{n=0}^{\infty} t^n\,U_n(x) && \text{(34.2-12b)}
\end{aligned}$$

Quick check of relation 34.2-12a using pari/gp:

```
? gen=truncate(taylor((1-t*x)/(1-2*x*t+t^2),t));
? for(k=0,5,print(k,":  ",polcoeff(gen,k,t)));
 0:  1
 1:  x
 2:  2*x^2 - 1
 3:  4*x^3 - 3*x
 4:  8*x^4 - 8*x^2 + 1
 5:  16*x^5 - 20*x^3 + 5*x
```

Binet forms for T (compare with relation 34.2-9) and U are

$$\begin{aligned}
T_n(z) &= \frac{1}{2}\left[\left(z+\sqrt{z^2-1}\right)^n + \left(z-\sqrt{z^2-1}\right)^n\right] && \text{(34.2-13a)}\\
U_n(z) &= \frac{1}{2\sqrt{z^2-1}}\left[\left(z+\sqrt{z^2-1}\right)^{n+1} - \left(z-\sqrt{z^2-1}\right)^{n+1}\right] && \text{(34.2-13b)}
\end{aligned}$$

Composition is multiplication of indices as can be seen by the definition (relation 34.2-1a on page 667):

$$T_n(T_m(x)) = T_{n\,m}(x) \tag{34.2-14}$$

For example,

$$\begin{aligned}
T_{2n}(x) &= T_2(T_n(x)) = 2\,T_n^2(x) - 1 && \text{(34.2-15a)}\\
&= T_n(T_2(x)) = T_n(2\,x^2-1) && \text{(34.2-15b)}
\end{aligned}$$

34.2.2 Index-doubling and relations between T and U

Index-doubling relations for the polynomials of the first kind are

$$\begin{aligned}
T_{2n} &= 2\,T_n^2 - 1 && \text{(34.2-16a)}\\
T_{2n+1} &= 2\,T_{n+1}\,T_n - x && \text{(34.2-16b)}\\
T_{2n-1} &= 2\,T_n\,T_{n-1} - x && \text{(34.2-16c)}
\end{aligned}$$

Similar relations for the polynomials of the second kind are

$$\begin{aligned}
U_{2n} &= U_n^2 - U_{n-1}^2 \;=\; (U_n + U_{n-1})\,(U_n - U_{n-1}) && \text{(34.2-17a)}\\
&= U_n\,(U_n - U_{n-2}) - 1 \;=\; U_{n-1}\,(U_{n+1} - U_{n-1}) + 1 && \text{(34.2-17b)}\\
U_{2n+1} &= U_n\,(U_{n+1} - U_{n-1}) && \text{(34.2-17c)}\\
&= 2\,U_n\,(U_{n+1} - x\,U_n) \;=\; 2\,U_n\,(x\,U_n - U_{n-1}) && \text{(34.2-17d)}\\
U_{2n-1} &= U_{n-1}\,(U_n - U_{n-2}) && \text{(34.2-17e)}\\
&= 2\,U_{n-1}\,(U_n - x\,U_{n-1}) \;=\; 2\,U_{n-1}\,(x\,U_{n-1} - U_{n-2}) && \text{(34.2-17f)}
\end{aligned}$$

Some relations between T and U are

$$\begin{aligned}
T_n &= U_n - x\,U_{n-1} \;=\; x\,U_{n-1} - U_{n-2} \;=\; \frac{1}{2}\,(U_n - U_{n-2}) && \text{(34.2-18a)}\\
T_{n+1} &= x\,T_n - (1-x^2)\,U_{n-1} && \text{(34.2-18b)}\\
U_{2n} &= 2\,T_n\,U_n - 1 && \text{(34.2-18c)}\\
U_{2n-1} &= 2\,T_n\,U_{n-1} \;=\; 2\,(T_{n+1}\,U_n + x) && \text{(34.2-18d)}\\
U_{2n+1} &= 2\,T_{n+1}\,U_n \;=\; 2\,(T_{n+2}\,U_{n-1} + x) && \text{(34.2-18e)}\\
U_{2^n-1} &= 2^n \prod_{k=0}^{n-1} T_{2^k} && \text{(34.2-18f)}
\end{aligned}$$

Relation 34.2-18b, written as

$$U_n \;=\; \frac{x\,T_{n+1} - T_{n+2}}{1-x^2} \;=\; \frac{T_n - x\,T_{n+1}}{1-x^2} \qquad \text{(34.2-19)}$$

can be used to compute the polynomials of the second kind from those of the first kind. One further has:

$$\begin{aligned}
T_{n+m} + T_{n-m} &= 2\,T_n\,T_m && \text{(34.2-20a)}\\
T_{n+m} - T_{n-m} &= 2\,(x^2+1)\,U_{n-1}\,U_{m-1} && \text{(34.2-20b)}\\
U_{n+m-1} + U_{n-m-1} &= 2\,U_{n-1}\,T_m && \text{(34.2-20c)}\\
U_{n+m-1} - U_{n-m-1} &= 2\,T_n\,U_{m-1} && \text{(34.2-20d)}
\end{aligned}$$

Expressions for certain sums:

$$\begin{aligned}
\sum_{k=0}^{n} T_{2k} &= \frac{1}{2}\,(1 + U_{2n}) && \text{(34.2-21a)}\\
\sum_{k=0}^{n-1} T_{2k+1} &= \frac{1}{2}\,U_{2n-1} && \text{(34.2-21b)}\\
\sum_{k=0}^{n} U_{2k} &= \frac{1 - T_{2n+2}}{2\,(1-x^2)} && \text{(34.2-21c)}\\
\sum_{k=0}^{n-1} U_{2k+1} &= \frac{x - T_{2n+1}}{2\,(1-x^2)} && \text{(34.2-21d)}
\end{aligned}$$

Using $\partial_x \cos(n\,\arccos(x)) = n\,\sin(n\,\arccos(x))/\sqrt{1-x^2}$ we obtain

$$\partial_x T_n(x) \;=\; n\,U_{n-1}(x) \qquad \text{(34.2-22)}$$

34.2.3 Fast computation of the Chebyshev polynomials

We give algorithms that improve on both the matrix power, and the polynomial based algorithms.

34.2.3.1 Chebyshev polynomials of the first kind

For even index use relation 34.2-16a ($T_{2n} = 2\,T_n^2 - 1$). For odd index we use relations 34.2-16c and 34.2-16b. We compute the pair $[T_{n-1},\, T_n]$ recursively via

$$[T_{n-1},\, T_n] \;=\; \left[2\,T_{q-1}\,T_q - x,\; 2\,T_q^2 - 1\right] \quad \text{where} \quad q = n/2, \quad \text{if } n \text{ even} \tag{34.2-23a}$$

$$[T_{n-1},\, T_n] \;=\; \left[2\,T_{q-1}^2 - 1,\; 2\,T_{q-1}\,T_q - x\right] \quad \text{where} \quad q = (n+1)/2, \quad \text{if } n \text{ odd} \tag{34.2-23b}$$

Note that no multiplication with x occurs thereby the computation is efficient also for floating point arguments. With integer x the cost of the computation of $T_n(x)$ is $\sim M(n)$ where $M(n)$ is the cost of a multiplication of numbers with the precision of the result. When x is a floating point number the cost is $\sim \log_2(n)\, M(n)$ where $M(n)$ is the cost of a multiplication with the precision used.

The code for the pair computations is

```
1   fvT(n, x)=
2   { /* return  [ T(n-1,x), T(n,x) ]  */
3       local(nr, t, t1, t2);
4       if ( n<=1,
5           if ( 1==n, return( [1, x] ) );
6           if ( 0==n, return( [x, 1] ) );
7           if ( -1==n, return( [2*x^2-1, x] ) );
8           return( 0 );  \\ disallow negative index < -1
9       );
10
11      nr = (n+1) >> 1; \\   if ( "n even",  nr = n/2 , nr = (n+1)/2; );
12      vr = fvT(nr, x);  \\ recursion
13      t1 = vr[1];  t2 = vr[2];
14      if ( !bitand(n,1), \\ n is even
15          t = [2*t1*t2-x, 2*t2^2-1];
16          ,
17          t = [2*t1^2-1, 2*t1*t2-x];
18       );
19      return( t );
20  }
```

The function to be called by the user is

```
1   fT(n, x)=
2   {
3       local(q, t, v, T);
4       n = abs(n);
5       if ( n<=1,
6           if (n>=0,  return(if(0==n,1,x)));
7           return( fT(-n, x) );
8       );
9       t = 0;  q = n;
10      while ( 0==bitand(q, 1), q>>=1; t+=1; );
11      \\ here: n==q*2^t
12      T = fvT(q, x)[2];
13      while ( t,  T=2*T*T-1; t-=1; );
14      return( T );
15  }
```

We check the speedup by comparing with the matrix-power computation that gives identical results. We compute $T_{4{,}545{,}967}(2)$, a number with more than 2,600,000 decimal digits:

```
vT(n,x)= return( ([1, x]*[0,-1; 1,2*x]^n) );
x=2;  \\ want integer calculations
n=4545967;
vT(n,x);   \\ computed in 9,800 ms.
fvT(n,x);  \\ computed in 2,241 ms.
```

C++ implementations for the computation of $T_n(2)$ and $T_n(x)$ modulo m are given in [FXT: mod/chebyshev1.cc].

34.2.3.2 Chebyshev polynomials of the second kind

One can use the fast algorithm for the polynomials of the first kind and relation 34.2-19 ($U_n = (T_n - x\,T_{n+1})/(1-x^2)$, involving a division):

```
fU(n, x)=
{
    local(v);
    if ( 1==x, return(n+1) );  \\ avoid division by zero
    v = fvT(n+1, x);
    return( (v[1]-x*v[2])/(1-x^2) );
}
```

We give an additional algorithm that uses 3 multiplication for each reduction of the index n. One multiplication is by the variable x. We compute the pair $[U_{n-1},\, U_n]$ recursively via

$$M_q := (U_q + U_{q-1})\,(U_q - U_{q-1}) \tag{34.2-24a}$$

$$[U_{n-1},\, U_n] = [2\,U_{q-1}\,(U_q - x\,U_{q-1})\,,\; M_q] \quad \text{where} \quad q = n/2, \quad \text{if } n \text{ even} \tag{34.2-24b}$$

$$[U_{n-1},\, U_n] = [M_q,\; 2\,U_q\,(x\,U_q - U_{q-1})] \quad \text{where} \quad q = (n-1)/2, \quad \text{if } n \text{ odd} \tag{34.2-24c}$$

The code for the pair computations is

```
fvU(n, x)=
{ /* return  [ U(n-1,x), U(n,x) ] */
    local(nr, u1, u0, ue, t, u);
    if ( n<=1,
        if ( 1==n, return( [1, (2*x)] ) );
        if ( 0==n, return( [0, 1] ) );
        if ( -1==n, return( [-1, 0] ) );
        if ( -2==n, return( [-(2*x), -1] ) );
        return( 0 );  \\ disallow negative index < -2
    );

    nr = n >> 1;  \\ if ( "n even",  nr = n/2 , nr = (n-1)/2; );
    vr = fvU(nr, x);  \\ recursion
    u1 = vr[1];  u0 = vr[2];
    ue = (u0+u1) * (u0-u1);
    if ( !bitand(n,1),  \\ n is even
        t = u1*(u0-x*u1);  t+=t;
        u = [t,  ue];
        ,
        t = u0*(x*u0-u1);  t+=t;
        u = [ue, t];
    );
    return( u );
}
```

The function to be called by the user is

```
fU(n, x)=  return( fvU(n,x)[2] );
```

The comparison with the matrix-power computation shows almost the same speedup as for the polynomials of the first kind:

```
vU(n,x)= return( [0, 1]*[0,-1; 1,2*x]^n );
x=2;  \\ want integer calculations
n=4545967;
vU(n,x);   \\ computed in 9,783 ms.
fvU(n,x);  \\ computed in 2,704 ms.
```

C++ implementations for the computation of $U_n(2)$ and $U_n(x)$ modulo m are given in [FXT: mod/chebyshev2.cc].

34.2.3.3 Symbolic computation

For symbolic computations the explicit power series as in 34.2-4a or 34.2-5a on page 669 should be preferred. The following routine computes T_n as a polynomial in x:

```
1   chebyTrec(n)=
2   {
3       local(b, s);
4       b = 2^(n-1);
5       if ( 0==n%2, if ( 0==n%4, s=+1, s=-1 ), s=0 );
6       forstep (k=n, 1, -2,
7           s += b*'x^(k);
8           b *= -(k*(k-1))/((n+k-2)*(n-k+2));
9       );
10      return( Pol(s) );
11  }
```

For U_n one can use

```
1   chebyUrec(n)=
2   {
3       local(b, s);
4       n += 1;
5       b = 2^(n-1)/n;
6       s = 0;
7       forstep (k=n, 1, -2,
8           s += (k)*b*'x^(k-1);
9           b *= -(k*(k-1))/((n+k-2)*(n-k+2));
10      );
11      return( Pol(s) );
12  }
```

34.2.4 Relations to approximations of the square root *

34.2.4.1 Padé approximants for $\sqrt{x^2 \pm 1}$

We start with the relation (from the definitions 34.2-1a and 34.2-1b on page 667, and $sin^2 + cos^2 = 1$)

$$T_n^2 - (x^2 - 1)\, U_{n-1}^2 \quad = \quad 1 \tag{34.2-25}$$

which we write as

$$\sqrt{x^2 - 1} \quad = \quad \sqrt{\frac{T_n^2 - 1}{U_{n-1}^2}} \tag{34.2-26}$$

If we define $R_n = T_n / U_{n-1}$, then

$$R_n(x) \quad = \quad \frac{T_n}{U_{n-1}} \;\approx\; \sqrt{x^2 - 1} \tag{34.2-27}$$

A composition law holds for R:

$$R_{m\,n}(x) \quad = \quad R_m(R_n(x)) \tag{34.2-28}$$

We list the first few values of $R_k(2)$ and $R_k(x)$:

k:	$R_k(2)$	$R_k(x)$
1:	$2/1$	$\frac{x}{1}$
2:	$7/4$	$\frac{2x^2 - 1}{2x}$
3:	$26/15$	$\frac{4x^3 - 3x}{4x^2 - 1}$
4:	$97/56$	$\frac{8x^4 - 8x^2 + 1}{8x^3 - 4x}$
5:	$362/209$	$\frac{16x^5 - 20x^3 + 5x}{16x^4 - 12x^2 + 1}$
∞:	$\sqrt{3}$	$\sqrt{x^2 - 1}$

(34.2-29)

If we define $T_n^+(x) := T(i\,x)/i^n$ and $U_n^+(x) := U(i\,x)/i^n$ then

$$T_n^{+\,2} - (x^2+1)\,U_{n-1}^{+\,2} \quad = \quad 1 \tag{34.2-30}$$

Defining $R_n^+ := T_n^+/U_{n-1}^+$ we have

$$R_{m\,n}^+(x) \quad = \quad R_m^+(R_n^+(x)) \tag{34.2-31}$$

and

$$\sqrt{x^2+1} \quad = \quad \sqrt{\frac{T_n^{+\,2}-1}{U_{n-1}^{+\,2}}} \quad \approx \quad \frac{T_n^+}{U_{n-1}^+} \;=\; R_n^+(x) \tag{34.2-32}$$

The first few values of $R_k^+(1)$ and $R_k^+(x)$ are

$$\begin{array}{cll} k\text{:} & R_k^+(1) & R_k^+(x) \\ 1\text{:} & 1/1 & \dfrac{x}{1} \\ 2\text{:} & 3/2 & \dfrac{2x^2+1}{2x} \\ 3\text{:} & 7/5 & \dfrac{4x^3+3x}{4x^2+1} \\ 4\text{:} & 17/12 & \dfrac{8x^4+8x^2+1}{8x^3+4x} \\ 5\text{:} & 41/29 & \dfrac{16x^5+20x^3+5x}{16x^4+12x^2+1} \\ \infty\text{:} & \sqrt{2} & \sqrt{x^2+1} \end{array} \tag{34.2-33}$$

Relations 34.2-30 and 34.2-25 can be used to power solutions of Pell's Diophantine equation, see section 37.13.2 on page 803.

34.2.4.2 Two products for the square root

For those fond of products: for $d > 0$, $d \neq 1$

$$\sqrt{d} \quad = \quad \prod_{k=0}^{\infty}\left(1+\frac{1}{q_k}\right) \quad \text{where} \quad q_0 = \frac{d+1}{d-1}, \quad q_{k+1} = 2\,q_k^2 - 1 \tag{34.2-34}$$

(convergence is quadratic) and

$$\sqrt{d} \quad = \quad \prod_{k=0}^{\infty}\left(1+\frac{2}{h_k}\right) \quad \text{where} \quad h_0 = \frac{d+3}{d-1}, \quad h_{k+1} = (h_k+2)^2\,(h_k-1)+1 \tag{34.2-35}$$

(convergence is cubic). These are given in [30] (also in [117], more expressions can be found in [109]). The paper gives $h_{k+1} = \frac{4\,d}{d-1}\prod_{i=0}^{k}\left(h_i^2\right) - 3$. Note that for relation 34.2-34 we have

$$q_k \quad = \quad T_{2^k}(q_0) \tag{34.2-36}$$

$$\frac{1}{q_k} \quad = \quad \frac{(d-1)^N}{\sum_{i=0}^{N}\binom{2\,N}{2\,i}\,d^i} \;=\; \frac{2\,(1-d)^N}{(1+\sqrt{d})^{2N}+(1-\sqrt{d})^{2N}} \quad \text{where} \quad N = 2^k \tag{34.2-37}$$

where T_n is the n-th Chebyshev polynomial of the first kind. One finds

$$q_k \quad = \quad T_{2^k}(1/c) \quad \text{where} \quad c = \frac{1-d}{1+d}, \quad c < 1 \tag{34.2-38}$$

and

$$\sqrt{\frac{1-c}{1+c}} \quad \approx \quad \frac{1-c}{c}\,\frac{U_{2^k-1}(1/c)}{T_{2^k}(1/c)} \tag{34.2-39}$$

which can be expressed in $d = \frac{1-c}{1+c}$ as

$$\sqrt{d} \quad \approx \quad \frac{2\,d}{1-d}\,\frac{U_{2^k-1}(\frac{1+d}{1-d})}{T_{2^k}(\frac{1+d}{1-d})} \quad \text{where} \quad d > 1 \tag{34.2-40}$$

where U_n is the n-th Chebyshev polynomial of the second kind. We have $U_{2^k-1}(x) = 2^k \prod_{i=0}^{k-1} T_{2^i}(x)$. Successively compute $T_{2^i} = 2\,T_{2^{i-1}}^2 - 1$ and accumulate the product $U_{2^i-1} = 2\,U_{2^{i-1}-1}\,T_{2^{i-1}}$ until U_{2^k-1} and T_{2^k} are obtained. Alternatively use the relation $U_k(x) = \frac{1}{k+1}\,\partial_x T_{k+1}(x)$ and use the recursion for the coefficients of T as shown on page 673.

A systematic approach to find product expressions for roots is given in section 28.7 on page 577.

Chapter 35

Cyclotomic polynomials, Hypergeometric functions, and continued fractions

We describe the cyclotomic polynomials and some of their properties, together with the Möbius inversion principle. We also give algorithms to convert a power series into Lambert series and infinite products.

We describe the hypergeometric functions which contain most of the 'useful' functions such as the logarithm and the sine as special cases. The transformation formulas for hypergeometric functions can be used to obtain series transformations that are non-obvious. The computation of certain hypergeometric functions by AGM-type algorithms is described in section 30.3 on page 598.

Further continued fractions are described together with algorithms for their computation.

35.1 Cylotomic polynomials, Möbius inversion, Lambert series

35.1.1 Cyclotomic polynomials

The roots (over $\mathbb{C}$) of the polynomial $x^n - 1$ are the n-th roots of unity:

$$x^n - 1 \quad = \quad \prod_{k=0}^{n-1} \left(x - \exp\left(\frac{2\,i\,\pi\,k}{n} \right) \right) \tag{35.1-1}$$

The n-th *cyclotomic polynomial* Y_n can be defined as the polynomial whose roots are the primitive n-th roots of unity:

$$Y_n(x) \quad := \quad \prod_{\substack{k=0..n-1 \\ \gcd(k,n)=1}} \left(x - \exp\left(\frac{2\,i\,\pi\,k}{n} \right) \right) \tag{35.1-2}$$

The degree of Y_n equals the number of primitive n-th roots, that is

$$\deg(Y_n) \quad = \quad \varphi(n) \tag{35.1-3}$$

The coefficients are integers, for example,

$$Y_{63}(x) \quad = \quad x^{36} - x^{33} + x^{27} - x^{24} + x^{18} - x^{12} + x^9 - x^3 + 1 \tag{35.1-4}$$

```
  n:   Yn(x)
  1:   x - 1
  2:   x + 1
  3:   x^2 + x + 1
  4:   x^2 + 1
  5:   x^4 + x^3 + x^2 + x + 1
  6:   x^2 - x + 1
  7:   x^6 + x^5 + x^4 + x^3 + x^2 + x + 1
  8:   x^4 + 1
  9:   x^6 + x^3 + 1
 10:   x^4 - x^3 + x^2 - x + 1
 11:   x^10 + ... + 1  <--= all coefficients are one for prime n
 12:   x^4 - x^2 + 1
 13:   x^12 + ... + 1
 14:   x^6 - x^5 + x^4 - x^3 + x^2 - x + 1
 15:   x^8 - x^7 + x^5 - x^4 + x^3 - x + 1
 16:   x^8 + 1
 17:   x^16 + ... + 1
 18:   x^6 - x^3 + 1
 19:   x^18 + ... + 1
 20:   x^8 - x^6 + x^4 - x^2 + 1
 21:   x^12 - x^11 + x^9 - x^8 + x^6 - x^4 + x^3 - x + 1
 22:   x^10 - x^9 + x^8 - x^7 + x^6 - x^5 + x^4 - x^3 + x^2 - x + 1
 23:   x^22 + ... + 1
 24:   x^8 - x^4 + 1
 25:   x^20 + x^15 + x^10 + x^5 + 1
 26:   x^12 - x^11 + x^10 - x^9 + x^8 - x^7 + x^6 - x^5 + x^4 - x^3 + x^2 - x + 1
 27:   x^18 + x^9 + 1
 28:   x^12 - x^10 + x^8 - x^6 + x^4 - x^2 + 1
 29:   x^28 + ... + 1
 30:   x^8 + x^7 - x^5 - x^4 - x^3 + x + 1
```

Figure 35.1-A: The first 30 cyclotomic polynomials.

The first 30 cyclotomic polynomials are shown in figure 35.1-A. The first cyclotomic polynomial with a coefficient not in the set $\{0, \pm 1\}$ is Y_{105}:

$$Y_{105}(x) \quad = \quad x^{48} + x^{47} + x^{46} - x^{43} - x^{42} - 2 \cdot x^{41} - x^{40} - x^{39} + \ldots \tag{35.1-5}$$

The cyclotomic polynomials are irreducible over $\mathbb{Z}$. All except Y_1 are self-reciprocal.

For n prime the cyclotomic polynomial $Y_n(x)$ equals $(x^n - 1)/(x - 1) = x^{n-1} + x^{n-2} + \ldots + x + 1$. For $n = 2\,k$ and odd $k \geq 3$ we have $Y_n(x) = Y_k(-x)$. For $n = p\,k$ where p is a prime that does not divide k we have $Y_n(x) = Y_k(x^p)/Y_p(x)$. The following algorithm for the computation of $Y_n(x)$ is given in [127, p.403]:

1. Let $[p_1, p_2, \ldots, p_r]$ the distinct prime divisors of n. Set $y_0(x) = x - 1$.

2. For $j = 1, 2, \ldots, r$ set $y_j(x) = y_j(x^{p_j})/y_j(x)$ (the division is exact).

3. Return $y_r(x^{n/(p_1\,p_2\,\cdots\,p_r)})$

The last statement uses the fact that for $n = k\,t$ where all prime factors of k divide t we have $Y_n(x) = Y_t(x^k)$. An implementation is

```
1  polcyclo2(n, z='x)=
2  {
3      local(fc, y);
4      fc = factor(n)[,1];  \\ prime divisors
5      y = z - 1;
6      for (j=1, #fc,  y=subst(y,z,z^fc[j])\y;  n\=fc[j]; );
7      y = subst(y, z, z^n);
8      return( y );
9  }
```

Note that the routine will only work when the argument z is a symbol.

35.1.2 The Möbius inversion principle

The *Möbius function* $\mu(n)$ is defined for positive integer arguments n as

$$\mu(n) \quad := \quad \begin{cases} 0 & \text{if} \quad n \quad \text{has a square factor} \\ (-1)^k & \text{if} \quad n \quad \text{is a product of } k \text{ distinct primes} \\ +1 & \text{if} \quad n = 1 \end{cases} \tag{35.1-6}$$

The function satisfies

$$\sum_{d \backslash n} \mu(d) \quad = \quad \begin{cases} 1 & \text{if} \quad n = 0 \\ 0 & \text{else} \end{cases} \tag{35.1-7}$$

A function $f(n)$ defined that satisfies

$$f(n \cdot m) \quad = \quad f(n) \cdot f(n) \quad \text{if} \quad \gcd(n, m) = 1 \tag{35.1-8}$$

is said to be *multiplicative*. For a multiplicative function one always has $f(1) = 1$ and $f(n) = f(p_1^{e_1}) \cdot f(p_2^{e_2}) \cdot \ldots \cdot f(p_k^{e_k})$ where $n = p_1^{e_1} \cdot p_2^{e_2} \cdot \ldots \cdot p_k^{e_k}$ is the factorization of n into distinct primes p_i. If the equality holds also for $\gcd(n, m) \neq 1$ the function is said to be *completely multiplicative*. Such a function satisfies $f(n) = f(p_1)^{e_1} \cdot f(p_2)^{e_2} \cdot \ldots \cdot f(p_k)^{e_k}$.

The Möbius function is multiplicative

$$\mu(n)\ \mu(m) \quad = \quad \begin{cases} \mu(n \cdot m) & \text{if} \quad \gcd(n, m) = 1 \\ 0 & \text{else} \end{cases} \tag{35.1-9}$$

For the cyclotomic polynomials one has

$$x^n - 1 \quad = \quad \prod_{d \backslash n} Y_d(x) \tag{35.1-10}$$

and

$$Y_n(x) \quad = \quad \prod_{d \backslash n} \left(x^d - 1\right)^{\mu(n/d)} \tag{35.1-11}$$

The relation implies a reasonably efficient algorithm for the computation of the cyclotomic polynomials. The method also works when the argument x is not a symbol.

n : $\mu(n)$	n : $\mu(n)$	n : $\mu(n)$	n : $\mu(n)$	n : $\mu(n)$	n : $\mu(n)$	n : $\mu(n)$	n : $\mu(n)$
1: +1	11: −1	21: +1	31: −1	41: −1	51: +1	61: −1	71: −1
2: −1	12: 0	22: +1	32: 0	42: −1	52: 0	62: +1	72: 0
3: −1	13: −1	23: −1	33: +1	43: −1	53: −1	63: 0	73: −1
4: 0	14: +1	24: 0	34: +1	44: 0	54: 0	64: 0	74: 1
5: −1	15: +1	25: 0	35: +1	45: 0	55: +1	65: +1	75: 0
6: +1	16: 0	26: +1	36: 0	46: +1	56: 0	66: −1	76: 0
7: −1	17: −1	27: 0	37: −1	47: −1	57: +1	67: −1	77: 1
8: 0	18: 0	28: 0	38: +1	48: 0	58: +1	68: 0	78: −1
9: 0	19: −1	29: −1	39: +1	49: 0	59: −1	69: +1	79: −1
10: +1	20: 0	30: −1	40: 0	50: 0	60: 0	70: −1	80: 0

Figure 35.1-B: Values of the Möbius function $\mu(n)$ for $n \leq 80$.

The pair of relations 35.1-10 and 35.1-11 is actually a special case of the (multiplicative) *Möbius inversion principle*:

$$g(n) \ = \ \prod_{d \backslash n} f(d) \quad \Longleftrightarrow \quad f(n) \ = \ \sum_{d \backslash n} g\,(d)^{\mu(n/d)} \tag{35.1-12}$$

Relation 35.1-10 implies (considering the polynomial degrees only and using relation 35.1-3)

$$n = \sum_{d\backslash n} \varphi(d) \tag{35.1-13}$$

while relation 35.1-11 corresponds to the equality

$$\varphi(n) = \sum_{d\backslash n} d\ \mu(n/d) \tag{35.1-14}$$

Relations 35.1-13 and 35.1-14 are a special case of the additive version of the Möbius inversion principle:

$$g(n) = \sum_{d\backslash n} f(d) \quad \Longleftrightarrow \quad f(n) = \sum_{d\backslash n} g(d)\ \mu(n/d) \tag{35.1-15}$$

More general, if $h(a\,b) = h(a)\,h(b)$ (see [103, p.447]) then

$$g(n) = \sum_{d\backslash n} f(d)\,h(n/d) \quad \Longleftrightarrow \quad f(n) = \sum_{d\backslash n} g(d)\,h(n/d)\ \mu(n/d) \tag{35.1-16}$$

Setting $h(n) = 1$ gives relation 35.1-15. The Möbius inversion principle is nicely explained in [133]. The sequence of the values of the Möbius function (see figure 35.1-B) is entry A008683 of [240].

We note two relations valid for multiplicative functions f:

$$\sum_{d\backslash n} \mu(d)\,f(d) = \prod_{d\backslash n,\ d \text{ prime}} (1 - f(d)) \tag{35.1-17a}$$

$$\sum_{d\backslash n} \mu(d)^2\,f(d) = \prod_{d\backslash n,\ d \text{ prime}} (1 + f(d)) \tag{35.1-17b}$$

Relation 35.1-17a with $f(n) = 1/n$ gives relation 35.1-13 and also

$$\varphi(n) = n \prod_{d\backslash n,\ d \text{ prime}} \left(1 - \frac{1}{d}\right) \tag{35.1-18}$$

35.1.3 Lambert series

A *Lambert series* is a series of the form

$$L(x) = \sum_{k>0} \frac{a_k\,x^k}{1 - x^k} = \sum_{k>0} \sum_{j>0} a_k\,x^{kj} \tag{35.1-19}$$

It can be converted to a Taylor series

$$L(x) = \sum_{k>0} b_k\,x^k \quad \text{where} \quad b_k = \sum_{d\backslash k} a_d \tag{35.1-20}$$

The inversion principle allows us to transform a Taylor series to a Lambert series:

$$a_k = \sum_{d\backslash k} b_d\ \mu(k/d) \tag{35.1-21}$$

With pari/gp the conversion to a Lambert series can be implemented as

```
ser2lambert(t)=
{
/* Let t=[a1,a2,a3, ...], n=length(v), where t(x)=sum_{k=1}^{n}{a_k*x^k};
 * Return  L=[l1,l2,l3,...] so that (up to order n)
 * t(x)=\sum_{j=1}^{n}{l_j*x^j/(1-x^j)}
 */
   local(n, L);
   n = length(t);
   L = vector(n);
   for (k=1, n,  fordiv(k, d, L[k]+=moebius(k/d)*t[d]); );
   return( L );
}
```

The conversion in the other direction is

```
lambert2ser(L)=
{ /* inverse of ser2lambert() */
   local(n, t);
   n = length(L);
   t = sum(k=1, length(L), O('x^(n+1))+L[k]*'x^k/(1-'x^k) );
   t = Vec(t);
   return( t );
}
```

We note a relation that is a useful for the computation of the sum, it is given in [169, p.644, ex.27],

$$L(x) \quad = \quad \sum_{k>0} x^{k^2} \left[a_k + \sum_{j>0} (a_k + a_{k+j})\, x^{k\,j} \right] \tag{35.1-22}$$

The special case $a_k = 1$ is given as relation 37.3-2 on page 761. For the related series

$$P(x) \quad = \quad \sum_{k>0} \frac{a_k\, x^k}{1+x^k} \;=\; -\sum_{k>0} \sum_{j>0} (-1)^j\, a_k\, x^{k\,j} \tag{35.1-23}$$

we find

$$P(x) \quad = \quad L(x) - 2\,L(x^2) \tag{35.1-24}$$

To verify the relation compute the k-th term on both sides: $a_k\, x^k/(1+x^k) = a_k\, x^k/(1-x^k) - 2\, a_k\, x^{2k}/(1-x^{2k})$. The other direction is obtained by repeatedly using $L(x) = P(x) + 2\,L(x^2)$:

$$L(x) \quad = \quad \sum_{k=0}^{\infty} 2^k\, P(x^{2^k}) \tag{35.1-25}$$

Use relations 35.1-22 and 35.1-24 to obtain

$$P(x) \quad = \quad \sum_{k>0} x^{k^2} \left[a_k\,(1 - 2\,x^{k^2}) + \sum_{j>0} (a_k + a_{k+j}) \left(x^{k\,j} - 2\,x^{(k+j)^2 - j^2} \right) \right] \tag{35.1-26}$$

35.1.4 Conversion of series to infinite products

Given a series with constant term one,

$$f(x) \quad = \quad 1 + \sum_{k>0} a_k\, x^k \tag{35.1-27}$$

we want to find an infinite product such that

$$f(x) \quad = \quad \prod_{k>0} (1 - x^k)^{b_k} \tag{35.1-28}$$

We take the logarithm, differentiate, and multiply by x:

$$x\,\frac{f'(x)}{f(x)} \quad = \quad \sum_{k>0} \frac{(-k\,b_k)\,x^k}{1-x^k} \tag{35.1-29}$$

The expression on the right hand side is a Lambert series with coefficients $-k\,b_k$, the expression on the left is easily computable as a power series, and we know how to compute a Lambert series from a power series. Thereby

$$b_k \quad = \quad -\frac{1}{k}\sum_{d\backslash k} q_k\;\mu(k/d) \tag{35.1-30}$$

where the q_k are the coefficients of the power series for $q(x) := x\,f'(x)/f(x)$. With pari/gp the conversion to a product can be implemented as

```
1   ser2prod(t)=
2   {
3   /* Let t=[1,a1,a2,a3, ...], n=length(v), where t(x)=1+sum_{k=1}^{n}{a_k*x^k};
4    * Return  p=[p1,p2,p3,...] so that (up to order n)
5    * t(x)=\prod_{j=1}^{n}{(1-x^j)^{p_j}}
6    */
7       local(v);
8       v = Ser(t);
9       v = v'/v;
10      v = vector(#t-1, j, polcoeff(v, j-1));
11      v = ser2lambert(v);
12      v = vector(#v, j, -v[j]/j);
13      return( v );
14  }
```

A simple example is $f(x) = \exp(x)$, so $x\,f'/f = x$, and

$$\exp(x) \quad = \quad \prod_{k>0} (1-x^k)^{-\,\mu(k)/k} \quad = \quad \frac{(1-x^2)^{1/2}\,(1-x^3)^{1/3}\,(1-x^5)^{1/5}\,\ldots}{(1-x^1)^{1/1}\,(1-x^6)^{1/6}\,(1-x^{10})^{1/10}\,\ldots} \tag{35.1-31}$$

Taking the logarithm, we obtain

$$x \quad = \quad -\sum_{k>0} \frac{\mu(k)}{k}\,\log(1-x^k) \tag{35.1-32}$$

Setting $f(x) = 1-2\,x$ we obtain relation 17.2-6a on page 362 (number of binary Lyndon words):

```
? ser2prod(Vec(1-2*x+O(x^20)))
  [2, 1, 2, 3, 6, 9, 18, 30, 56, 99, 186, 335, 630, 1161, 2182, 4080, 7710, 14532, 27594]
```

Setting $f(x) = 1-x-x^2$ gives the number of binary Lyndon words without the subsequence 00 (entry A006206 of [240]):

```
? ser2prod(Vec(1-x-x^2+O(x^20)))
  [1, 1, 1, 1, 2, 2, 4, 5, 8, 11, 18, 25, 40, 58, 90, 135, 210, 316, 492]
```

The ordinary generating function for the e_k corresponding to the product form $f(x) = \prod (1-x^k)^{e_k}$ is

$$\sum_{k=1}^{\infty} e_k\,x^k \quad = \quad -\sum_{k=1}^{\infty} \frac{\mu(k)}{k}\,\log\big(f(x^k)\big) \tag{35.1-33}$$

This can be seen by using the product form for f on the right hand side, using the power series $\log(1-x) = -(x+x^2/2+x^3/3+\ldots)$, and using the defining property of the Möbius function (relation 35.1-7 on page 679). An example is relation 17.2-6b on page 362. For the cyclotomic polynomials we obtain (via relation 35.1-11 on page 679):

$$-\sum_{k=1}^{\infty} \frac{\mu(k)}{k}\,\log\big(Y_n(x^k)\big) \quad = \quad \sum_{d\backslash n} \mu(d)\,x^{n/d} \tag{35.1-34}$$

For example, setting $n = 2$ we obtain

$$x^2 - x \quad = \quad -\sum_{k=1}^{\infty} \frac{\mu(k)}{k} \log\left(1 + x^k\right) \tag{35.1-35}$$

35.1.4.1 An alternative product form

For the transformation into products of the form $\prod (1 + x^k)^{c_k}$ we set

$$f(x) \quad = \quad \prod_{k>0} (1 + x^k)^{c_k} \tag{35.1-36}$$

and note that

$$x\,\frac{f'(x)}{f(x)} \quad = \quad \sum_{k>0} \frac{(+k\,c_k)\,x^k}{1 + x^k} \tag{35.1-37}$$

So we need a transformation into series of this type. As the Möbius transform is not (easily) applicable we use a greedy algorithm:

```
1   ser2lambertplus(t)=
2   {
3   /* Let t=[a1,a2,a3, ...], n=length(v), where t(x)=sum_{k=1}^{n}{a_k*x^k};
4    * Return  L=[l1,l2,l3,...] so that (up to order n)
5    * t(x)=\sum_{j=1}^{n}{l_j*x^j/(1+x^j)}
6    */
7       local(n, L, k4);
8       n = length(t);
9       L = vector(n);
10      for (k=1, n,
11          tk = t[k];
12          L[k] = tk;
13          \\ subtract tk * x^k/(1+x^k):
14          forstep(j=k, n, 2*k, t[j] -= tk);
15          forstep(j=k+k, n, 2*k, t[j] += tk);
16      );
17      return( L );
18  }
```

Now we can compute the product form via

```
1   ser2prodplus(t)=
2   {
3   /* Let t=[1,a1,a2,a3, ...], n=length(v), where t(x)=1+sum_{k=1}^{n}{a_k*x^k};
4    * Return  p=[p1,p2,p3,...] so that (up to order n)
5    * t(x)=\prod_{j=1}^{n}{(1+x^j)^{p_j}}
6    */
7       local(v);
8       v = Ser(t);
9       v = v'/v;
10      v = vector(#t-1, j, polcoeff(v, j-1));
11      v = ser2lambertplus(v);
12      v = vector(#v, j, v[j]/j);
13      return( v );
14  }
```

A product $\prod_{k>0} (1 - x^k)^{b_k}$ can be converted into a product $\prod_{k>0} (1 + x^k)^{c_k}$ via the relation $(1 - x) = \prod_{k\geq 0} (1 + x^{2^k})^{-1}$.

35.1.4.2 Conversion to eta-products

The conversion of a series to a product of the form (*eta-product*, or η-*product*),

$$\prod_{k=1}^{\infty} \left[\eta(x^k)\right]^{u_k} \quad \text{where} \quad \eta(x) := \prod_{j=1}^{\infty} 1 - x^j \tag{35.1-38}$$

```
   r= 2:  ( E(y^2)^3 ) / ( E(y^4) )
   r= 3:  ( E(y^3)^4 ) / ( E(y^9) )
   r= 4:  ( E(y^4)^7 ) / ( E(y^8)^3 )
   r= 5:  ( E(y^5)^6 ) / ( E(y^25) )
   r= 6:  ( E(y^6)^12 E(y^36) ) / ( E(y^12)^4 E(y^18)^3 )
   r= 7:  ( E(y^7)^8 ) / ( E(y^49) )
   r= 8:  ( E(y^8)^15 ) / ( E(y^16)^7 )
   r= 9:  ( E(y^9)^13 ) / ( E(y^27)^4 )
   r=10:  ( E(y^10)^18 E(y^100) ) / ( E(y^20)^6 E(y^50)^3 )
   r=11:  ( E(y^11)^12 ) / ( E(y^121) )
   r=12:  ( E(y^12)^28 E(y^72)^3 ) / ( E(y^24)^12 E(y^36)^7 )
   r=13:  ( E(y^13)^14 ) / ( E(y^169) )
   r=14:  ( E(y^14)^24 E(y^196) ) / ( E(y^28)^8 E(y^98)^3 )
   r=15:  ( E(y^15)^24 E(y^225) ) / ( E(y^45)^6 E(y^75)^4 )
   r=16:  ( E(y^16)^31 ) / ( E(y^32)^15 )
   r=17:  ( E(y^17)^18 ) / ( E(y^289) )
   r=18:  ( E(y^18)^39 E(y^108)^4 ) / ( E(y^36)^13 E(y^54)^12 )
   r=19:  ( E(y^19)^20 ) / ( E(y^361) )
   r=20:  ( E(y^20)^42 E(y^200)^3 ) / ( E(y^40)^18 E(y^100)^7 )
   r=21:  ( E(y^21)^32 E(y^441) ) / ( E(y^63)^8 E(y^147)^4 )
   r=22:  ( E(y^22)^36 E(y^484) ) / ( E(y^44)^12 E(y^242)^3 )
   r=23:  ( E(y^23)^24 ) / ( E(y^529) )
   r=24:  ( E(y^24)^60 E(y^144)^7 ) / ( E(y^48)^28 E(y^72)^15 )
   r=25:  ( E(y^25)^31 ) / ( E(y^125)^6 )
   r=26:  ( E(y^26)^42 E(y^676) ) / ( E(y^52)^14 E(y^338)^3 )
   r=27:  ( E(y^27)^40 ) / ( E(y^81)^13 )
   r=28:  ( E(y^28)^56 E(y^392)^3 ) / ( E(y^56)^24 E(y^196)^7 )
   r=29:  ( E(y^29)^30 ) / ( E(y^841) )
   r=30:  ( E(y^30)^72 E(y^180)^6 E(y^300)^4 E(y^450)^3 ) /
            ( E(y^60)^24 E(y^90)^18 E(y^150)^12 E(y^900) )
   r=31:  ( E(y^31)^32 ) / ( E(y^961) )
   r=32:  ( E(y^32)^63 ) / ( E(y^64)^31 )
   r=33:  ( E(y^33)^48 E(y^1089) ) / ( E(y^99)^12 E(y^363)^4 )
```

Figure 35.1-C: Functions $\eta_r(y) := \prod_{j=0}^{r-1} \eta(\omega^j\, y)$ as products of η-functions.

can be done by a greedy algorithm:

```
1   etaprod(v)=
2   {
3   /* Let t=[1,a1,a2,a3, ...], n=length(v), where t(x)=1+sum_{k=1}^{n}{a_k*x^k};
4    * Return  p=[p1,p2,p3,...] so that (up to order n)
5    * t(x)=\prod_{j=1}^{n}{eta(x^j)^{p_j}}
6    * where eta(x) = prod(k>0, (1-x^k))
7    */
8       local(n, t);
9       v = ser2prod(v);
10      n = length(v);
11      for (k=1, n,
12          t = v[k];
13          forstep (j=k+k, n, k, v[j]-=t; );
14      );
15      return( v );
16  }
```

Similarly, to convert into a product of the form

$$\prod_{k=1}^{\infty} \left[\eta_+(x^k)\right]^{u_k} \quad \text{where} \quad \eta_+(x) := \prod_{j=1}^{\infty} 1 + x^j \tag{35.1-39}$$

use

```
1   etaprodplus(v)=
2   {
3   /* Let t=[1,a1,a2,a3, ...], n=length(v), where t(x)=1+sum_{k=1}^{n}{a_k*x^k};
4    * Return  p=[p1,p2,p3,...] so that (up to order n)
5    * t(x)=\prod_{j=1}^{n}{eta_+(x^j)^{p_j}}
6    * where eta_+(x) = prod(k>0, (1+x^k))
```

```
 7    */
 8       local(n, t);
 9       v = ser2prodplus(v);
10       n = length(v);
11       for (k=1, n,
12           t = v[k];
13           forstep (j=k+k, n, k, v[j]-=t; );
14       );
15       return( v );
16   }
```

The routines are useful for computations with the generating functions of partitions of certain types, see section 14.4 on page 332. Here we just give:

$$\eta(-x) = \frac{\eta(x^2)^3}{\eta(x)\,\eta(x^4)} \tag{35.1-40a}$$

$$\eta(+i\,x)\,\eta(-i\,x) = \frac{\eta(x^4)^8}{\eta(x^2)^3\,\eta(x^8)^3} \tag{35.1-40b}$$

Figure 35.1-C gives more product formulas.

35.2 Hypergeometric functions

The *hypergeometric function* $F\left(\begin{matrix}a,\,b\\c\end{matrix}\middle|\,z\right)$ can be defined as

$$F\left(\begin{matrix}a,\,b\\c\end{matrix}\middle|\,z\right) := \sum_{k=0}^{\infty} \frac{a^{\overline{k}}\,b^{\overline{k}}}{c^{\overline{k}}}\,\frac{z^k}{k!} \tag{35.2-1}$$

where $z^{\overline{k}} := z\,(z+1)\,(z+2)\,\ldots\,(z+k-1)$ is the *rising factorial power* ($z^{\overline{0}} := 1$). Some sources use the so-called *Pochhammer symbol* $(x)_k$ which is the same: $(x)_k = x^{\overline{k}}$. We'll stick to the factorial notation.

The variable z is the *argument* of the function, a, b and c are the *parameters*. Parameters in the upper and lower row are called upper and lower parameters, respectively.

Note the $k! = 1^{\overline{k}}$ in the denominator of relation 35.2-1. You might want to have the hidden lower parameter 1 in mind:

$$F\left(\begin{matrix}2,\,2\\1\end{matrix}\middle|\,z\right) = \text{“}F\left(\begin{matrix}2,\,2\\1,\,1\end{matrix}\middle|\,z\right)\text{”} \tag{35.2-2}$$

The expression is a sum of perfect squares if z is a square.

We have

$$F\left(\begin{matrix}a,\,b\\c\end{matrix}\middle|\,z\right) = 1 + \frac{a}{1}\,\frac{b}{c}\,z\,\left(1 + \frac{a+1}{2}\,\frac{b+1}{c+1}\,z\,\left(1 + \frac{a+2}{3}\,\frac{b+2}{c+2}\,z\,(1+\ldots)\right)\right) \tag{35.2-3}$$

so by formula 35.2-3 hypergeometric functions with rational arguments can be computed with the binary splitting method described in section 32.1.

Hypergeometric functions can have any number of parameters:

$$F\left(\begin{matrix}a_1,\,\ldots,\,a_m\\b_1,\,\ldots,\,b_n\end{matrix}\middle|\,z\right) = \sum_{k=0}^{\infty} \frac{a_1^{\overline{k}}\ldots a_m^{\overline{k}}}{b_1^{\overline{k}}\ldots b_n^{\overline{k}}}\,\frac{z^k}{k!} \tag{35.2-4}$$

These are sometimes called *generalized hypergeometric functions*. The number of upper and lower parameters are often emphasized as subscripts left and right to the symbol F. For example, ${}_mF_n$ for the hypergeometric function in the last relation.

The functions $F\left(\begin{matrix}a\\b\end{matrix}\middle|\, z\right)$ (of type ${}_1F_1$) are sometimes written as $M(a,b,z)$ or $\Phi(a;b;z)$. *Kummer's function* $U(a,b,z)$ (or $\Psi(a;b;z)$) is related to hypergeometric functions of type ${}_2F_0$:

$$U(a,b,z) \;=\; z^{-a}\,F\left(\begin{matrix}a,\,1+a-b\\ \end{matrix}\middle|\, -1/z\right) \tag{35.2-5}$$

Note that series ${}_2F_0$ are not convergent. Still, they can be used as asymptotic series for large values of z.

The so-called *Whittaker functions* are related to hypergeometric functions as follows:

$$M_{a,b}(z) \;=\; e^{-z/2}\,z^{b+1/2}\,F\left(\begin{matrix}\frac{1}{2}+b-a\\1+2b\end{matrix}\middle|\, z\right) \tag{35.2-6a}$$

$$W_{a,b}(z) \;=\; e^{-z/2}\,z^{b+1/2}\,U\left(\frac{1}{2}+b-a,\,1+2b,\,z\right) \tag{35.2-6b}$$

$$\;=\; e^{-z/2}\,z^{a}\,F\left(\begin{matrix}\frac{1}{2}+b-a,\,\frac{1}{2}-b-a\\ \end{matrix}\middle|\, -1/z\right) \tag{35.2-6c}$$

Negative integer parameters in the upper row lead to polynomials:

$$F\left(\begin{matrix}-3,\,3\\1\end{matrix}\middle|\, z\right) \;=\; 1-9\,z+18\,z^2-10\,z^3 \tag{35.2-7}$$

The lower parameter must not be zero or a negative integer unless there is a negative upper parameter with smaller absolute value.

Sometimes one finds the notational convention to omit an argument $z=1$:

$$F\left(\begin{matrix}a_1,\,\ldots,\,a_m\\b_1,\,\ldots,\,b_n\end{matrix}\right) \;:=\; F\left(\begin{matrix}a_1,\,\ldots,\,a_m\\b_1,\,\ldots,\,b_n\end{matrix}\middle|\, 1\right) \tag{35.2-8}$$

In what follows the argument is never omitted.

An in-depth treatment of hypergeometric functions is [15].

35.2.1 Derivative and differential equation

Using the relation

$$\frac{d}{dz^n}\,F\left(\begin{matrix}a,\,b\,\ldots\\c,\,\ldots\end{matrix}\middle|\, z\right) \;=\; \frac{a^{\overline{n}}\,b^{\overline{n}}\,\ldots}{c^{\overline{n}}\,\ldots}\,F\left(\begin{matrix}a+n,\,b+n\,\ldots\\c+n,\,\ldots\end{matrix}\middle|\, z\right) \tag{35.2-9}$$

one can verify that $f(z)=F\left(\begin{matrix}a,b\\c\end{matrix}\middle|\, z\right)$ is a solution of the differential equation

$$z\,(1-z)\,\frac{d^2f}{dz^2}+[c-(1+a+b)\,z]\,\frac{df}{dz}-a\,b\,f \;=\; 0 \tag{35.2-10}$$

A general form of the differential equation satisfied by $F\left(\begin{matrix}a,b,c,\ldots\\u,v,w,\ldots\end{matrix}\middle|\, z\right)$ is

$$z\,(\vartheta+a)\,(\vartheta+b)\,(\vartheta+c)\,\ldots\,f(z) \;=\; \vartheta\,(\vartheta+u-1)\,(\vartheta+v-1)\,(\vartheta+w-1)\,\ldots\,f(z) \tag{35.2-11}$$

where ϑ is the operator $z\,\frac{d}{dz}$. The leftmost ϑ on the right hand side of the equation takes care of the hidden lower parameter 1: $\vartheta=(\vartheta+1-1)$. See [133] for a beautiful derivation. Relation 10.15-4a on page 282 can be used to rewrite powers of ϑ as polynomials in $\frac{d}{dz}$.

35.2.2 Evaluations for fixed z

A closed form (in terms of the gamma function) evaluation at $z = 1$ can be given for ${}_2F_1$:

$$F\left(\begin{matrix}a,\, b\\ c\end{matrix}\middle|\, 1\right) \quad = \quad \frac{\Gamma(c)\,\Gamma(c-a-b)}{\Gamma(c-a)\,\Gamma(c-b)} \qquad \text{if} \quad \mathfrak{Re}(c-a-b) > 0 \quad \text{or } b \in \mathbb{N}, b < 0 \tag{35.2-12}$$

When $c - a - b < 0$ then [273, ex.18, p.299]

$$\lim_{z\to 1-} F\left(\begin{matrix}a,\, b\\ c\end{matrix}\middle|\, z\right) \Big/ \left(\frac{\Gamma(c)\,\Gamma(a+b-c)}{\Gamma(a)\,\Gamma(b)}\,(1-z)^{c-a-b}\right) \quad = \quad 1 \tag{35.2-13a}$$

and, for $c - a - b = 0$,

$$\lim_{z\to 1-} F\left(\begin{matrix}a,\, b\\ c\end{matrix}\middle|\, z\right) \Big/ \left(\frac{\Gamma(a+b)}{\Gamma(a)\,\Gamma(b)}\,\log\frac{1}{1-z}\right) \quad = \quad 1 \tag{35.2-13b}$$

For $z = -1$ there is an evaluation due to Kummer:

$$F\left(\begin{matrix}a,\, b\\ 1+a-b\end{matrix}\middle|\, -1\right) \quad = \quad \frac{\Gamma(1-a+b)\,\Gamma(1+a/2)}{\Gamma(1+a)\,\Gamma(1+a/2-b)} \tag{35.2-14a}$$

$$\quad = \quad 2^{-a}\,\pi\,\frac{\Gamma(1-a+b)}{\Gamma(1/2+a/2)\,\Gamma(1+a/2-b)} \tag{35.2-14b}$$

Several evaluations at $z = \frac{1}{2}$ are given in [1], we just give one:

$$F\left(\begin{matrix}a,\, b\\ \frac{1}{2}+\frac{1}{2}a+\frac{1}{2}b\end{matrix}\middle|\, \frac{1}{2}\right) \quad = \quad \sqrt{\pi}\,\frac{\Gamma(\frac{1}{2}+\frac{1}{2}a+\frac{1}{2}b)}{\Gamma(\frac{1}{2}+\frac{1}{2}a)\,\Gamma(\frac{1}{2}+\frac{1}{2}b)} \tag{35.2-15}$$

For further information see (chapter 15 of) [1], [268] and [214]. Various evaluations of $F\left(\begin{smallmatrix}-an,\, bn+b_1\\ cn+c_1\end{smallmatrix}\middle|\, z\right)$ for integer a, b, c and n, $1 \le a \le 2$, $-4 \le b \le 4$ and $-4 \le c \le 4$ can be found in [114].

35.2.3 Extraction of even and odd part

Let $E[f(z)] = (f(z)+f(z))/2$ (the even powers of the series of $f(z)$), and $O[f(z)] = (f(z)-f(-z))/2$ (the odd powers). We express the even and odd parts of a hypergeometric series as hypergeometric functions:

$$E\left[F\left(\begin{matrix}a,\, b\\ c\end{matrix}\middle|\, z\right)\right] \quad = \quad F\left(\begin{matrix}\frac{a}{2},\, \frac{a+1}{2},\, \frac{b}{2},\, \frac{b+1}{2}\\ \frac{c}{2},\, \frac{c+1}{2},\, \frac{1}{2}\end{matrix}\middle|\, z^2\right) \tag{35.2-16a}$$

$$O\left[F\left(\begin{matrix}a,\, b\\ c\end{matrix}\middle|\, z\right)\right] \quad = \quad \frac{a\,b}{c}\,z\,F\left(\begin{matrix}\frac{a+1}{2},\, \frac{a+2}{2},\, \frac{b+1}{2},\, \frac{b+2}{2}\\ \frac{c+1}{2},\, \frac{c+2}{2},\, \frac{3}{2}\end{matrix}\middle|\, z^2\right) \tag{35.2-16b}$$

The lowers parameters 1/2 and 3/2 are due to the hidden lower parameter 1. The general case for

$$H(z) \quad := \quad F\left(\begin{matrix}a_1,\, \ldots,\, a_m\\ b_1,\, \ldots,\, b_n\end{matrix}\middle|\, z\right) \tag{35.2-17a}$$

is

$$E\left[H(z)\right] \quad = \quad F\left(\begin{matrix}\frac{a_1}{2},\, \frac{a_1+1}{2},\, \ldots,\, \frac{a_m}{2},\, \frac{a_m+1}{2}\\ \frac{b_1}{2},\, \frac{b_1+1}{2},\, \ldots,\, \frac{b_n}{2},\, \frac{b_n+1}{2},\, \frac{1}{2}\end{matrix}\middle|\, X\,z^2\right) \tag{35.2-17b}$$

$$O\left[H(z)\right] \quad = \quad \frac{a_1\,\cdots\,a_m}{b_1\,\cdots\,b_n}\,z\,F\left(\begin{matrix}\frac{a_1+1}{2},\, \frac{a_1+2}{2},\, \ldots,\, \frac{a_m+1}{2},\, \frac{a_m+2}{2}\\ \frac{b_1+1}{2},\, \frac{b_1+2}{2},\, \ldots,\, \frac{b_n+1}{2},\, \frac{b_n+2}{2},\, \frac{3}{2}\end{matrix}\middle|\, X\,z^2\right) \tag{35.2-17c}$$

where $X = 4^{m-n-1}$. For example,

$$E\left[F\left(\begin{matrix}1\\2\end{matrix}\middle| z\right)\right] \;=\; F\left(\begin{matrix}\frac{1}{2},\,1\\1,\,\frac{3}{2},\,\frac{1}{2}\end{matrix}\middle|\,\frac{z^2}{4}\right) \;=\; F\left(\begin{matrix}\\\frac{3}{2}\end{matrix}\middle|\,\frac{z^2}{4}\right) \;=\; \frac{\sinh z}{z} \tag{35.2-18}$$

We indicate a further generalization by the extraction of all terms of $H(z)$ where the exponent of z is divisible by 3:

$$\frac{H(z)+H(\omega\,z)+H(\omega^2\,z)}{3} \;=\; F\left(\begin{matrix}\frac{a_1}{3},\,\frac{a_1+1}{3},\,\frac{a_1+2}{3}\,\ldots,\,\frac{a_m}{3},\,\frac{a_m+1}{3},\,\frac{a_m+2}{3}\\\frac{b_1}{3},\,\frac{b_1+1}{3},\,\frac{b_1+2}{3},\,\ldots,\,\frac{b_n}{3},\,\frac{b_n+1}{3},\,\frac{b_n+2}{3},\,\frac{1}{3},\,\frac{2}{3}\end{matrix}\middle|\,X\,z^3\right) \tag{35.2-19}$$

where $\omega = \exp(2\,i\,\pi/3)$ and $X = 27^{m-n-1}$. For example, with $H(z) = \exp(z) = F\left(\begin{matrix}\\\end{matrix}\middle| z\right)$ we obtain

$$F\left(\begin{matrix}\\\frac{1}{3},\,\frac{2}{3}\end{matrix}\middle|\,\frac{z^3}{27}\right) \;=\; \sum_{k=0}^{\infty}\frac{z^{3k}}{(3k)!} \tag{35.2-20}$$

Define the power series $C_j(z)$, for $j \in \{0,1,2\}$, by

$$C_s(z) \;=\; \sum_{k=0}^{\infty}\frac{z^{3k+s}}{(3k+s)!} \tag{35.2-21a}$$

then (omitting arguments)

$$\det\begin{bmatrix}C_0 & C_1 & C_2\\ C_2 & C_0 & C_1\\ C_1 & C_2 & C_0\end{bmatrix} \;=\; C_0^3 + C_1^3 + C_2^3 - 3\,C_0\,C_1\,C_2 \;=\; 1 \tag{35.2-21b}$$

which is a three power series analogy to the relation $\cosh^2 - \sinh^2 = 1$.

For the extraction of the coefficient at the positions j mod M replace every upper and lower parameter A by the M parameters $(A+j)/M$, $(A+j+1)/M$, $(A+j+2)/M$, …, $(A+j+M-1)/M$, and the argument z by $X\,z^M$ where $X = (M^M)^{m-n-1}$.

35.2.4 Transformations

As obvious from the definition, parameters in the upper row can be swapped (capitalized symbols for readability):

$$F\left(\begin{matrix}A,\,B,\,c\\e,\,f,\,g\end{matrix}\middle| z\right) \;=\; F\left(\begin{matrix}B,\,A,\,c\\e,\,f,\,g\end{matrix}\middle| z\right) \tag{35.2-22}$$

The same is true for the lower row. Usually one writes the parameters in ascending order. Identical elements in the lower and upper row can be canceled:

$$F\left(\begin{matrix}a,\,b,\,C\\e,\,f,\,C\end{matrix}\middle| z\right) \;=\; F\left(\begin{matrix}a,\,b\\e,\,f\end{matrix}\middle| z\right) \tag{35.2-23}$$

These trivial transformations are true for any number of elements. The following transformations are only valid for the given structure, unless the list of parameters contain an ellipsis '…'.

35.2.4.1 Elementary relations

$$F\left(\begin{matrix}a,\,b,\,\ldots\\ c,\,\ldots\end{matrix}\Big|\,z\right) = 1 + z\,\frac{a\,b\,\ldots}{c\,\ldots}\,F\left(\begin{matrix}a+1,\,b+1,\,\ldots,\,1\\ c+1,\,\ldots,\,2\end{matrix}\Big|\,z\right) \tag{35.2-24}$$

$$(a-b)\,F\left(\begin{matrix}a,\,b,\,\ldots\\ c,\,\ldots\end{matrix}\Big|\,z\right) = a\,F\left(\begin{matrix}a+1,\,b,\,\ldots\\ c,\,\ldots\end{matrix}\Big|\,z\right) - b\,F\left(\begin{matrix}a,\,b+1,\,\ldots\\ c,\,\ldots\end{matrix}\Big|\,z\right) \tag{35.2-25}$$

$$(a-c)\,F\left(\begin{matrix}a,\,b,\,\ldots\\ c+1,\,\ldots\end{matrix}\Big|\,z\right) = a\,F\left(\begin{matrix}a+1,\,b,\,\ldots\\ c+1,\,\ldots\end{matrix}\Big|\,z\right) - c\,F\left(\begin{matrix}a,\,b,\,\ldots\\ c,\,\ldots\end{matrix}\Big|\,z\right) \tag{35.2-26}$$

These are given in [133], the following is taken from [273].

$$F\left(\begin{matrix}a,\,b\\ c\end{matrix}\Big|\,z\right) = F\left(\begin{matrix}a,\,b+1\\ c\end{matrix}\Big|\,z\right) - \frac{a\,z}{c}\,F\left(\begin{matrix}a+1,\,b+1\\ c+1\end{matrix}\Big|\,z\right) \tag{35.2-27}$$

More relations of this type are given in [1].

35.2.4.2 Pfaff's reflection law and Euler's identity

Pfaff's reflection law can be given as either of:

$$\frac{1}{(1-z)^a}\,F\left(\begin{matrix}a,\,b\\ c\end{matrix}\Big|\,\frac{-z}{1-z}\right) = F\left(\begin{matrix}a,\,c-b\\ c\end{matrix}\Big|\,z\right) \tag{35.2-28a}$$

$$F\left(\begin{matrix}a,\,b\\ c\end{matrix}\Big|\,z\right) = \frac{1}{(1-z)^a}\,F\left(\begin{matrix}a,\,c-b\\ c\end{matrix}\Big|\,\frac{-z}{1-z}\right) \tag{35.2-28b}$$

$$= \frac{1}{(1-z)^b}\,F\left(\begin{matrix}c-a,\,b\\ c\end{matrix}\Big|\,\frac{-z}{1-z}\right) \tag{35.2-28c}$$

Euler's identity is obtained by applying the Pfaff reflection on both upper parameters:

$$F\left(\begin{matrix}a,\,b\\ c\end{matrix}\Big|\,z\right) = (1-z)^{(c-a-b)}\,F\left(\begin{matrix}c-a,\,c-b\\ c\end{matrix}\Big|\,z\right) \tag{35.2-29}$$

Now write Euler's transform as

$$F\left(\begin{matrix}a,\,b\\ a+b+r\end{matrix}\Big|\,z\right)/F\left(\begin{matrix}a+r,\,b+r\\ a+b+r\end{matrix}\Big|\,z\right) = (1-z)^r \tag{35.2-30}$$

If both the numerator and the denominator terminate then the expression on the left is a Padé approximant for the $1/r$-th root, see section 28.2.3 on page 563.

Euler's transformation can generalized for hypergeometric functions ${}_{r+1}F_r$, see [190]. We give two transforms for hypergeometrics ${}_3F_2$ where one upper parameter exceeds a lower parameter by one, taken from [192, p.17]. The first one is reminiscent to the Pfaff reflection, we have

$$F\left(\begin{matrix}a_1,\,a_2,\,e+1\\ b,\,e\end{matrix}\Big|\,z\right) = (1-z)^{-a_1}\,F\left(\begin{matrix}a_1,\,b-a_2-1,\,f+1\\ b,\,f\end{matrix}\Big|\,\frac{-z}{1-z}\right) \tag{35.2-31}$$

where $f = e\,\frac{b-a_2-1}{e-a_2}$. The second one is similar to Euler's identity:

$$F\left(\begin{matrix}a_1,\,a_2,\,e+1\\ b,\,e\end{matrix}\Big|\,z\right) = (1-z)^{b-a_1-a_2-1}\,F\left(\begin{matrix}b-a_1-1,\,b-a_2-1,\,g+1\\ b,\,g\end{matrix}\Big|\,z\right) \tag{35.2-32}$$

where $g = \frac{(b-a_1-1)\,(b-a_2-1)\,e}{(b-a_1-a_2-1)\,e+a_1\,a_2}$.

35.2.4.3 A transformation by Gauss

$$F\left(\begin{matrix}2a,\,2b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,z\right) = F\left(\begin{matrix}a,\,b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,4z(1-z)\right) \quad \text{where} \quad |z|<\frac{1}{2} \tag{35.2-33a}$$

$$F\left(\begin{matrix}a,\,b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,z\right) = F\left(\begin{matrix}2a,\,2b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,\frac{1-\sqrt{1-z}}{2}\right) \tag{35.2-33b}$$

Note that the right hand side of relation 35.2-33a does not change if z is replaced by $1-z$, so it seems that

$$F\left(\begin{matrix}2a,\,2b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,z\right) = F\left(\begin{matrix}2a,\,2b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,1-z\right) \tag{35.2-34}$$

However, the relation is true only for terminating series, that is, for polynomials. Rewriting relation 35.2-33a for the argument $\frac{1-z}{2}$ we obtain

$$F\left(\begin{matrix}2a,\,2b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,\frac{1-z}{2}\right) = F\left(\begin{matrix}a,\,b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,1-z^2\right) \tag{35.2-35}$$

35.2.4.4 Whipple's identity and quadratic transformations

Whipple's identity connects two hypergeometric functions ${}_3F_2$:

$$F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}a+\frac{1}{2},\,1-a-b-c\\ 1+a-b,\,1+a-c\end{matrix}\Big|\,\frac{-4z}{(1-z)^2}\right) = (1-z)^a\,F\left(\begin{matrix}a,\,b,\,c\\ 1+a-b,\,1+a-c\end{matrix}\Big|\,z\right) \tag{35.2-36}$$

Specializing 35.2-36 for $c=(a+1)/2$ (note the symmetry between b and c so specializing for $c=(b+1)/2$ produces the identical relation) gives

$$F\left(\begin{matrix}a,\,b\\ 1+a-b\end{matrix}\Big|\,z\right) = \frac{1}{(1-z)^a}\,F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}a+\frac{1}{2}-b\\ 1+a-b\end{matrix}\Big|\,\frac{-4z}{(1-z)^2}\right) \tag{35.2-37}$$

$$F\left(\begin{matrix}a,\,b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,z\right) = \left(\frac{2\,(1-\sqrt{1-z})}{z}\right)^{2a} F\left(\begin{matrix}2a,\,a-b+\frac{1}{2}\\ a+b+\frac{1}{2}\end{matrix}\Big|\,-\frac{\left(1-\sqrt{1-z}\right)^2}{z}\right) \tag{35.2-38}$$

With $c:=a-b$ in 35.2-37 one obtains:

$$F\left(\begin{matrix}a,\,a-c\\ 1+c\end{matrix}\Big|\,z\right) = \frac{1}{(1-z)^a}\,F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}-\frac{1}{2}a+c\\ 1+c\end{matrix}\Big|\,\frac{-4z}{(1-z)^2}\right) \tag{35.2-39}$$

Similarly as for the relations by Gauss, from relations 35.2-37 and 35.2-38:

$$F\left(\begin{matrix}a,\,b\\ 1+a-b\end{matrix}\Big|\,-\frac{1-z}{1+z}\right) = \left(\frac{1+z}{2}\right)^a F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}a+\frac{1}{2}-b\\ 1+a-b\end{matrix}\Big|\,1-z^2\right) \tag{35.2-40a}$$

$$F\left(\begin{matrix}a,\,b\\ 1+a-b\end{matrix}\Big|\,-\frac{1-\sqrt{1-z^2}}{1+\sqrt{1-z^2}}\right) = \left(\frac{1+\sqrt{1-z^2}}{2}\right)^a F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}a+\frac{1}{2}-b\\ 1+a-b\end{matrix}\Big|\,z^2\right) \tag{35.2-40b}$$

Relations 35.2-40b and 35.2-40a can be obtained from each other by setting $x=\sqrt{1-y^2}$ (and replacing y by x). The same is true for the next pair of relations:

$$F\left(\begin{matrix}a,\,b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,1-z^2\right) = \left(\frac{2}{1+z}\right)^{2a} F\left(\begin{matrix}2a,\,a-b+\frac{1}{2}\\ a+b+\frac{1}{2}\end{matrix}\Big|\,-\frac{1-z}{1+z}\right) \tag{35.2-41a}$$

$$F\left(\begin{matrix}a,\,b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,z^2\right) = \left(\frac{2}{1+\sqrt{1-z^2}}\right)^{2a} F\left(\begin{matrix}2a,\,a-b+\frac{1}{2}\\ a+b+\frac{1}{2}\end{matrix}\Big|\,-\frac{1-\sqrt{1-z^2}}{1+\sqrt{1-z^2}}\right) \tag{35.2-41b}$$

The transformations

$$F\left(\begin{matrix}a,\,b\\ a-b+1\end{matrix}\Big|\,z\right) = (1+z)^{-a}\,F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}a+\frac{1}{2}\\ a-b+1\end{matrix}\Big|\,\frac{4z}{(1+z)^2}\right) \quad (35.2\text{-}42a)$$

$$= (1-z)^{-a}\,F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}a-b+\frac{1}{2}\\ a-b+1\end{matrix}\Big|\,\frac{-4z}{(1-z)^2}\right) \quad (35.2\text{-}42b)$$

$$= (1\pm\sqrt{z})^{-2a}\,F\left(\begin{matrix}a,\,a-b+\frac{1}{2}\\ 2a-2b+1\end{matrix}\Big|\,\frac{\pm 4\sqrt{z}}{(1\pm\sqrt{z})^2}\right) \quad (35.2\text{-}42c)$$

are given in [1]. Specializing for $a=b$ gives

$$F\left(\begin{matrix}a,\,a\\ 1\end{matrix}\Big|\,z\right) = (1+z)^{-a}\,F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}+\frac{1}{2}a\\ 1\end{matrix}\Big|\,\frac{4z}{(1+z)^2}\right) \quad (35.2\text{-}42d)$$

$$= (1-z)^{-a}\,F\left(\begin{matrix}\frac{1}{2}a,\,\frac{1}{2}-\frac{1}{2}a\\ 1\end{matrix}\Big|\,\frac{-4z}{(1-z)^2}\right) \quad (35.2\text{-}42e)$$

$$= (1\pm\sqrt{z})^{-2a}\,F\left(\begin{matrix}a,\,\frac{1}{2}\\ 1\end{matrix}\Big|\,\frac{\pm 4\sqrt{z}}{(1\pm\sqrt{z})^2}\right) \quad (35.2\text{-}42f)$$

Relation 35.2-42e can be obtained by setting $c=0$ in relation 35.2-37. Observe that the hypergeometric function on the right hand side of relation 35.2-42e does not change when replacing a by $1-a$. The next $({}_3F_2)$ transformation is given in [190]:

$$(1-z)^{-1}\,F\left(\begin{matrix}a,\,b,\,1\\ \frac{1}{2}+\frac{1}{2}a+\frac{1}{2}b,\,2\end{matrix}\Big|\,z\right) = F\left(\begin{matrix}\frac{1}{2}+\frac{1}{2}a,\,\frac{1}{2}+\frac{1}{2}b,\,1\\ \frac{1}{2}+\frac{1}{2}a+\frac{1}{2}b,\,2\end{matrix}\Big|\,4z(1-z)\right) \quad (35.2\text{-}43)$$

The following are special cases of this transformation:

$$(1-z)^{-1}\,F\left(\begin{matrix}a,\,1-a\\ 2\end{matrix}\Big|\,z\right) = F\left(\begin{matrix}\frac{1}{2}+\frac{1}{2}a,\,1-\frac{1}{2}a\\ 2\end{matrix}\Big|\,4z(1-z)\right) \quad (35.2\text{-}44a)$$

$$(1-z)^{-1}\,F\left(\begin{matrix}a,\,1\\ 2\end{matrix}\Big|\,z\right) = F\left(\begin{matrix}\frac{1}{2}+\frac{1}{2}a,\,1+\frac{1}{2}a,\,1\\ 1+a,\,2\end{matrix}\Big|\,4z(1-z)\right) \quad (35.2\text{-}44b)$$

$$(1-z)^{-1}\,F\left(\begin{matrix}a,\,1\\ \frac{3}{2}+\frac{1}{2}a\end{matrix}\Big|\,z\right) = F\left(\begin{matrix}\frac{1}{2}+\frac{1}{2}a,\,\frac{3}{2},\,1\\ \frac{3}{2}+\frac{1}{2}a,\,2\end{matrix}\Big|\,4z(1-z)\right) \quad (35.2\text{-}44c)$$

The nonlinear transformation (given in [214, p.21])

$$F\left(\begin{matrix}a,\,b\\ c\end{matrix}\Big|\,z\right) = (1-\omega)^{2a}\sum_{n=0}^{\infty} d_n\,\omega^n \quad (35.2\text{-}45a)$$

where

$$\omega = \frac{-4z}{(1-z)^2}, \qquad z = \frac{\sqrt{1-\omega}-1}{\sqrt{1-\omega}+1} \quad (35.2\text{-}45b)$$

and

$$d_0 = 1 \quad (35.2\text{-}45c)$$

$$d_1 = \frac{2a\,(c-2b)}{c} \quad (35.2\text{-}45d)$$

$$d_{n+2} = \frac{2\,(c-2b)\,(n+1+a)\,d_{n+1}+(n+2a)\,(n+2a+1-c)\,d_n}{(n+2)\,(n+1+c)} \quad (35.2\text{-}45e)$$

maps the complex (z-)plane into the unit circle. Thereby the ω-form of the series converges for all $z\neq 1$.

35.2.4.5 Clausen's product formulas

Clausen's formulas connect hypergeometric functions of type ${}_2F_1$ and ${}_3F_2$:

$$\left[F\left(\begin{matrix}a,\,b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,z\right)\right]^2 \;=\; F\left(\begin{matrix}2a,\,a+b,\,2b\\ a+b+\frac{1}{2},\,2a+2b\end{matrix}\Big|\,z\right) \tag{35.2-46a}$$

$$F\left(\begin{matrix}\frac{1}{4}+a,\,\frac{1}{4}+b\\ 1+a+b\end{matrix}\Big|\,z\right)\,F\left(\begin{matrix}\frac{1}{4}-a,\,\frac{1}{4}-b\\ 1-a-b\end{matrix}\Big|\,z\right) \;=\; F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}+a-b,\,\frac{1}{2}-a+b\\ 1+a+b,\,1-a-b\end{matrix}\Big|\,z\right) \tag{35.2-46b}$$

If $a = b + \frac{1}{2}$ in 35.2-46a then (two parameters on the right hand side cancel)

$$\left[F\left(\begin{matrix}b+\frac{1}{2},\,b\\ 2b+1\end{matrix}\Big|\,z\right)\right]^2 \;=\; F\left(\begin{matrix}2b+\frac{1}{2},\,2b\\ 4b+1\end{matrix}\Big|\,z\right) \tag{35.2-47}$$

and the right hand side again matches the structure on the left. The corresponding function can be identified (see [56, p.190]) as $G_b(z) := F\left(\begin{matrix}b+\frac{1}{2},\,b\\ 2b+1\end{matrix}\Big|\,z\right) = \left(\frac{1+\sqrt{1-z}}{2}\right)^{-2b}$. We have $G_{n\,m}(z) = [G_n(z)]^m$.

Specializing relation 35.2-46b for $b = -a$ we obtain

$$\left[F\left(\begin{matrix}\frac{1}{4}+a,\,\frac{1}{4}-a\\ 1\end{matrix}\Big|\,z\right)\right]^2 \;=\; F\left(\begin{matrix}\frac{1}{2}+2a,\,\frac{1}{2},\,\frac{1}{2}-2a\\ 1,\,1\end{matrix}\Big|\,z\right) \tag{35.2-48}$$

For $a = 0$, $z = 1$ (or z a sixth power) this relation is an identity between the square of a sum of squares and a sum of cubes:

$$\left[\sum_{n=0}^{\infty}\left[\prod_{j=1}^{n}\frac{1/4+j-1}{j}\right]^2\right]^2 \;=\; \sum_{n=0}^{\infty}\left[\prod_{j=1}^{n}\frac{1/2+j-1}{j}\right]^3 \;=\; 1.39320392968\ldots \tag{35.2-49}$$

The relation can be obtained by setting $\alpha = \beta = 1/4$ and $\gamma = 1/2$ in exercise 16 in [273, p.298]. Setting $a = 1/2$ in exercise 28 in [273, p.301] we find that the quantity equals $\pi/\Gamma(3/4)^4 = \Gamma(1/4)^4/(4\,\pi^3)$. For the square root of the expressions ($\sqrt{1.39320\ldots} = 1.180340\ldots$) we have [115, p.34]:

$$F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\ 1\end{matrix}\Big|\,\frac{1}{2}\right) \;=\; \left[\sum_{n=-\infty}^{\infty}e^{-n^2\,\pi}\right]^2 \;=\; 1.180340599016\ldots \tag{35.2-50}$$

We note that relation 35.2-46a can be obtained as the special case $c = a + b$ of the following relation given in [273, ex.16, p.298]:

$$F\left(\begin{matrix}a,\,b\\ c+\frac{1}{2}\end{matrix}\Big|\,z\right)\,F\left(\begin{matrix}c-a,\,c-b\\ c+\frac{1}{2}\end{matrix}\Big|\,z\right) \;=\; \sum_{k=0}^{\infty}A_k\,\frac{c^{\overline{k}}}{\left(c+\frac{1}{2}\right)^{\overline{k}}}\,z^k \tag{35.2-51a}$$

where the A_k are defined by

$$(1-z)^{a+b-c}\,F\left(\begin{matrix}2a,\,2b\\ 2c\end{matrix}\Big|\,z\right) \;=\; \sum_{k=0}^{\infty}A_k\,z^k \tag{35.2-51b}$$

The following relations are given in [15, p.184]:

$$F\left(\begin{matrix}a,\,b\\ a+b-\frac{1}{2}\end{matrix}\Big|\,z\right)\,F\left(\begin{matrix}a,\,b\\ a+b+\frac{1}{2}\end{matrix}\Big|\,z\right) \;=\; F\left(\begin{matrix}2a,\,2b,\,a+b\\ 2a+2b-1,\,a+b+\frac{1}{2}\end{matrix}\Big|\,z\right) \tag{35.2-52a}$$

$$F\left(\begin{matrix}a,\,b\\ a+b-\frac{1}{2}\end{matrix}\Big|\,z\right)\,F\left(\begin{matrix}a,\,b-1\\ a+b-\frac{1}{2}\end{matrix}\Big|\,z\right) \;=\; F\left(\begin{matrix}2a,\,2b-1,\,a+b-1\\ 2a+2b-2,\,a+b+\frac{1}{2}\end{matrix}\Big|\,z\right) \tag{35.2-52b}$$

35.2.4.6 The Kummer transformation

The Kummer transformation connects two hypergeometric functions of type ${}_1F_1$:

$$\exp(z)\,F\left(\begin{matrix}a\\a+b\end{matrix}\middle|\,-z\right) \;=\; F\left(\begin{matrix}b\\a+b\end{matrix}\middle|\,z\right) \tag{35.2-53}$$

The relation is not valid if both a and b are negative integers. In that case one obtains the Padé approximants of $\exp(z)$, see relation 31.2-17 on page 625.

A transformation from ${}_1F_1$ to ${}_2F_3$ is given by

$$F\left(\begin{matrix}a\\b\end{matrix}\middle|\,z\right)\,F\left(\begin{matrix}a\\b\end{matrix}\middle|\,-z\right) \;=\; F\left(\begin{matrix}a,\,b-a\\b,\,\frac{1}{2}b,\,\frac{1}{2}(b+1)\end{matrix}\middle|\,\frac{z^2}{4}\right) \tag{35.2-54}$$

Setting $b = 2a$ and using 35.2-53 gives

$$\left[F\left(\begin{matrix}a\\2a\end{matrix}\middle|\,z\right)\right]^2 \;=\; \exp(z)\,F\left(\begin{matrix}a\\a+\frac{1}{2},\,2a\end{matrix}\middle|\,\frac{z^2}{4}\right) \tag{35.2-55}$$

The following transformation connects functions ${}_0F_1$ and ${}_2F_3$:

$$F\left(\begin{matrix}\\a\end{matrix}\middle|\,z\right)\,F\left(\begin{matrix}\\b\end{matrix}\middle|\,z\right) \;=\; F\left(\begin{matrix}\frac{1}{2}(a+b),\,\frac{1}{2}(a+b-1)\\a,\,b,\,a+b-1\end{matrix}\middle|\,4z\right) \tag{35.2-56}$$

Setting $b = a$ gives (cancellation of parameters on the right hand side)

$$\left[F\left(\begin{matrix}\\a\end{matrix}\middle|\,z\right)\right]^2 \;=\; F\left(\begin{matrix}a-\frac{1}{2}\\a,\,2a-1\end{matrix}\middle|\,4z\right) \tag{35.2-57}$$

From relations 35.2-55 and 35.2-57 one can obtain

$$\exp(z)\,F\left(\begin{matrix}\\a\end{matrix}\middle|\,\frac{z^2}{4}\right) \;=\; F\left(\begin{matrix}a-\frac{1}{2}\\2a-1\end{matrix}\middle|\,2z\right) \tag{35.2-58}$$

The following relations can be derived from the preceding ones:

$$\left[F\left(\begin{matrix}\frac{1}{2}(a+b),\,\frac{1}{2}(a+b-1)\\a,\,b,\,a+b-1\end{matrix}\middle|\,z\right)\right]^2 \;=\; F\left(\begin{matrix}a-\frac{1}{2}\\a,\,2a-1\end{matrix}\middle|\,z\right)\,F\left(\begin{matrix}b-\frac{1}{2}\\b,\,2b-1\end{matrix}\middle|\,z\right) \tag{35.2-59}$$

$$F\left(\begin{matrix}a\\2a\end{matrix}\middle|\,z\right)\,F\left(\begin{matrix}a\\2a\end{matrix}\middle|\,-z\right) \;=\; \left[F\left(\begin{matrix}\\a+\frac{1}{2}\end{matrix}\middle|\,\frac{z^2}{16}\right)\right]^2 \tag{35.2-60a}$$

$$F\left(\begin{matrix}a\\2a+1\end{matrix}\middle|\,z\right)\,F\left(\begin{matrix}a\\2a+1\end{matrix}\middle|\,-z\right) \;=\; F\left(\begin{matrix}a\\2a+1,\,a+\frac{1}{2}\end{matrix}\middle|\,\frac{z^2}{4}\right) \tag{35.2-60b}$$

$$F\left(\begin{matrix}\\a\end{matrix}\middle|\,z\right)\,F\left(\begin{matrix}\\1-a\end{matrix}\middle|\,z\right) \;=\; \frac{1}{2}\left[1+F\left(\begin{matrix}\frac{1}{2}\\a,\,1-a\end{matrix}\middle|\,4z\right)\right] \tag{35.2-61a}$$

$$F\left(\begin{matrix}\\a\end{matrix}\middle|\,z\right)\,F\left(\begin{matrix}\\a+1\end{matrix}\middle|\,z\right) \;=\; F\left(\begin{matrix}a+\frac{1}{2}\\a+1,\,2a\end{matrix}\middle|\,4z\right) \tag{35.2-61b}$$

35.2.5 Examples: elementary functions

The 'well-known' functions like exp, log and sin are expressed as hypergeometric functions. In some cases a transformation is applied to give an alternative series.

35.2.5.1 Powers, roots, and binomial series

$$\frac{1}{(1-z)^a} = F\left(\begin{matrix}a\\ \end{matrix}\middle| z\right) = \sum_{k=0}^{\infty}\binom{a+k-1}{k}z^k = F\left(\begin{matrix}-a\\ \end{matrix}\middle| \frac{-z}{1-z}\right) \tag{35.2-62a}$$

$$(1+z)^a = F\left(\begin{matrix}-a\\ \end{matrix}\middle| -z\right) = \sum_{k=0}^{\infty}\binom{a}{k}z^k = F\left(\begin{matrix}a\\ \end{matrix}\middle| \frac{z}{1+z}\right) \tag{35.2-62b}$$

An important special case of relation 35.2-62a is

$$\frac{1}{1-z} = F\left(\begin{matrix}1\\ \end{matrix}\middle| z\right) = \sum_{k=0}^{\infty}z^k = F\left(\begin{matrix}\frac{1}{2},\,1\\ 2\end{matrix}\middle| 4z(1-z)\right) \quad \text{where} \quad z<\frac{1}{2} \tag{35.2-63}$$

The last equality is obtained by setting $a=0$ in relation 35.2-44a on page 691.

$$F\left(\begin{matrix}-s,\,s+1\\ s\end{matrix}\middle| z\right) = (1-2z)\,(1-z)^{s-1} \tag{35.2-64}$$

$$F\left(\begin{matrix}\frac{n}{2},\,\frac{n+1}{2}\\ \frac{1}{2}\end{matrix}\middle| z\right) = \frac{(1-\sqrt{z})^{-n}+(1+\sqrt{z})^{-n}}{2} \tag{35.2-65a}$$

$$F\left(\begin{matrix}\frac{n}{2},\,\frac{n+1}{2}\\ n+1\end{matrix}\middle| z\right) = \left(\frac{1+\sqrt{1-z}}{2}\right)^{-n} \tag{35.2-65b}$$

$$F\left(\begin{matrix}\frac{n}{2},\,\frac{n+1}{2}\\ n\end{matrix}\middle| z\right) = \frac{1}{\sqrt{1-z}}\left(\frac{2}{1+\sqrt{1-z}}\right)^{n-1} \tag{35.2-65c}$$

$$F\left(\begin{matrix}\frac{n+1}{2},\,\frac{n+2}{2}\\ \frac{3}{2}\end{matrix}\middle| z\right) = \frac{(1-\sqrt{z})^{-n}-(1+\sqrt{z})^{-n}}{2\,n\sqrt{z}} \quad \text{if} \quad n\neq 0 \tag{35.2-66a}$$

$$= \frac{1}{2\sqrt{z}}\log\left(\frac{1+\sqrt{z}}{1-\sqrt{z}}\right) \quad \text{if} \quad n=0 \tag{35.2-66b}$$

$$\frac{(1+z)^n-(1-z)^n}{(1+z)^n+(1-z)^n} = n\,z\,F\left(\begin{matrix}\frac{1}{2}-\frac{n}{2},\,1-\frac{n}{2}\\ \frac{3}{2}\end{matrix}\middle| z^2\right)\Big/F\left(\begin{matrix}\frac{1}{2}-\frac{n}{2},\,-\frac{n}{2}\\ \frac{1}{2}\end{matrix}\middle| z^2\right) \tag{35.2-67a}$$

$$= n\,z\,F\left(\begin{matrix}\frac{1}{2}-\frac{n}{2},\,\frac{1}{2}+\frac{n}{2}\\ \frac{3}{2}\end{matrix}\middle| \frac{z^2}{z^2-1}\right)\Big/F\left(\begin{matrix}\frac{1}{2}-\frac{n}{2},\,\frac{1}{2}+\frac{n}{2}\\ \frac{1}{2}\end{matrix}\middle| \frac{z^2}{z^2-1}\right) \tag{35.2-67b}$$

35.2.5.2 Chebyshev polynomials

The Chebyshev polynomials are treated in section 34.2 on page 667, we have

$$T_n(1-2z) = F\left(\begin{matrix}n,\,-n\\ \frac{1}{2}\end{matrix}\middle| z\right) \tag{35.2-68a}$$

$$T_n(z) = F\left(\begin{matrix}n,\,-n\\ \frac{1}{2}\end{matrix}\middle| \frac{1-z}{2}\right) \tag{35.2-68b}$$

$$U_n(z) = (n+1)\,F\left(\begin{matrix}-n,\,n+2\\ \frac{3}{2}\end{matrix}\middle| \frac{1-z}{2}\right) \tag{35.2-68c}$$

Using relation 34.2-14 on page 670 (as $T_n(T_{1/n}(z)) = z = \mathrm{id}(z)$) we find that

$$F\left(\begin{matrix}n,\,-n\\ \frac{1}{2}\end{matrix}\Big|\,\frac{1-z}{2}\right)^{[-1]} \;=\; F\left(\begin{matrix}\frac{1}{n},\,-\frac{1}{n}\\ \frac{1}{2}\end{matrix}\Big|\,\frac{1-z}{2}\right) \tag{35.2-69}$$

near $z = 1$ (here $F^{[-1]}$ denotes the inverse function).

35.2.5.3 Hermite polynomials

The Hermite polynomials $H_n(z)$ can be defined by the recurrence

$$H_{n+1}(z) \;=\; 2z\,H_n(z) - 2n\,H_{n-1}(x) \tag{35.2-70}$$

where $H_0(z) = 1$ and $H_1(z) = 2z$. The first few are

$$\begin{aligned}
H_0 &= 1 \qquad (35.2\text{-}71)\\
H_1 &= 2z\\
H_2 &= 4z^2 - 2\\
H_3 &= 8z^3 - 12z\\
H_4 &= 16z^4 - 48z^2 + 12\\
H_5 &= 32z^5 - 160z^3 + 120z\\
H_6 &= 64z^6 - 480z^4 + 720z^2 - 120\\
H_7 &= 128z^7 - 1344z^5 + 3360z^3 - 1680z\\
H_8 &= 256z^8 - 3584z^6 + 13440z^4 - 13440z^2 + 1680\\
H_9 &= 512z^9 - 9216z^7 + 48384z^5 - 80640z^3 + 30240z\\
H_{10} &= 1024z^{10} - 23040z^8 + 161280z^6 - 403200z^4 + 302400z^2 - 30240
\end{aligned}$$

For nonnegative integer n we have

$$H_n(z) \;=\; (2z)^n\,F\left(\begin{matrix}-\frac{1}{2}\,n,\,-\frac{1}{2}\,(n-1)\\ \end{matrix}\Big|\,-\frac{1}{z^2}\right) \tag{35.2-72}$$

35.2.5.4 Stirling numbers of the first kind

A generating hypergeometric function for the (unsigned) Stirling numbers of the first kind $s(n,m)$ (see section 10.15 on page 282) is given by

$$F\left(\begin{matrix}1,\,e\\ \end{matrix}\Big|\,z\right) \;=\; \sum_{n=0}^{\infty} e^{\overline{m}}\,z^n \;=\; \sum_{n=0}^{\infty}\left(\sum_{m=1}^{m} e^m\,s(n,m)\right)z^n \tag{35.2-73a}$$

$$\begin{aligned}
&= 1 + e\,z + \left(e + e^2\right)z^2 + \left(2e + 3e^2 + e^3\right)z^3 + \qquad (35.2\text{-}73b)\\
&\quad + \left(6e + 11e^2 + 6e^3 + e^4\right)z^4 +\\
&\quad + \left(24e + 50e^2 + 35e^3 + 10e^4 + e^5\right)z^5 +\\
&\quad + \left(120e + 274e^2 + 225e^3 + 85e^4 + 15e^5 + e^6\right)z^6 + \ldots
\end{aligned}$$

35.2.5.5 Logarithm and exponential function

$$\log(1+z) \;=\; z\,F\left(\begin{matrix}1,\,1\\ 2\end{matrix}\Big|\,-z\right) \;=\; \sum_{k=0}^{\infty}\frac{(-1)^k\,z^{k+1}}{k+1} \tag{35.2-74a}$$

$$\log\left(\frac{1+z}{1-z}\right) \;=\; 2\,z\,F\left(\begin{matrix}\frac{1}{2},\,1\\ \frac{3}{2}\end{matrix}\Big|\,z^2\right) \;=\; \log\left(1 + \frac{2}{1-z}\right) \tag{35.2-74b}$$

For large arguments z the following relation can be useful:

$$\log(1+z) \;=\; -\log\left(1-\frac{z}{1+z}\right) \;=\; \frac{z}{1+z}\,F\left(\left.\begin{matrix}1,\,1\\2\end{matrix}\right|\frac{z}{1+z}\right) \tag{35.2-75}$$

$$\exp(z) \;=\; F\left(\left.\begin{matrix}\\\end{matrix}\right|z\right) \;=\; \sum_{k=0}^{\infty}\frac{z^k}{k!} \tag{35.2-76}$$

35.2.5.6 Bessel functions and error function

The Bessel functions J_n of the first kind, and the modified Bessel functions I_n (as given in [1]):

$$J_n(z) \;=\; \frac{(z/2)^n}{n!}\,F\left(\left.\begin{matrix}\\n+1\end{matrix}\right|\frac{-z^2}{4}\right) \tag{35.2-77a}$$

$$I_n(z) \;=\; \frac{(z/2)^n}{n!}\,F\left(\left.\begin{matrix}\\n+1\end{matrix}\right|\frac{z^2}{4}\right) \tag{35.2-77b}$$

$$F\left(\left.\begin{matrix}\\n\end{matrix}\right|z\right) \;=\; \frac{(n-1)!}{z^{(n-1)/2}}\,I_{n-1}(2\sqrt{z}) \tag{35.2-77c}$$

$$F\left(\left.\begin{matrix}\\1\end{matrix}\right|z\right) \;=\; I_0(2\sqrt{z}) \;=\; \sum_{k=0}^{\infty}\frac{z^k}{k!^2} \tag{35.2-77d}$$

Error function (the Kummer transformation, relation 35.2-53 on page 693, gives relation 35.2-78b):

$$\frac{\sqrt{\pi}}{2}\,\mathrm{erf}(z) \;:=\; \int_{t=0}^{z} e^{-t^2}\,dt \;=\; z\,F\left(\left.\begin{matrix}\frac{1}{2}\\\frac{3}{2}\end{matrix}\right|-z^2\right) \tag{35.2-78a}$$

$$\;=\; z\,e^{-z^2}\,F\left(\left.\begin{matrix}1\\\frac{3}{2}\end{matrix}\right|z^2\right) \;=\; z\,e^{-z^2}\sum_{k=0}^{\infty}\frac{(2z^2)^k}{1\cdot 3\cdot 5\cdots(2k+1)} \tag{35.2-78b}$$

$$\;=\; \frac{1}{2\,x}\left[F\left(\left.\begin{matrix}-\frac{1}{2}\\\frac{1}{2}\end{matrix}\right|-z^2\right)-e^{-z^2}\right] \tag{35.2-78c}$$

35.2.5.7 Trigonometric and hyperbolic functions

Series for sine and cosine:

$$\sin(z) \;=\; z\,F\left(\left.\begin{matrix}\\\frac{3}{2}\end{matrix}\right|\frac{-z^2}{4}\right) \;=\; \sum_{k=0}^{\infty}\frac{(-1)^k\,z^{2k+1}}{(2k+1)!} \tag{35.2-79a}$$

$$\sinh(z) \;=\; z\,F\left(\left.\begin{matrix}\\\frac{3}{2}\end{matrix}\right|\frac{z^2}{4}\right) \;=\; \sum_{k=0}^{\infty}\frac{z^{2k+1}}{(2k+1)!} \tag{35.2-79b}$$

Applying the transformation 35.2-57 on page 693 to relation 35.2-79a gives

$$[\sin(z)]^2 \;=\; z^2\,F\left(\left.\begin{matrix}1\\\frac{3}{2},\,2\end{matrix}\right|-z^2\right) \tag{35.2-80}$$

$$\cos(z) \;=\; F\left(\left.\begin{matrix}\\\frac{1}{2}\end{matrix}\right|\frac{-z^2}{4}\right) \;=\; \sum_{k=0}^{\infty}\frac{(-1)^k\,z^{2k}}{(2k)!} \tag{35.2-81a}$$

$$\cosh(z) \;=\; F\left(\left.\begin{matrix}\\\frac{1}{2}\end{matrix}\right|\frac{z^2}{4}\right) \;=\; \sum_{k=0}^{\infty}\frac{z^{2k}}{(2k)!} \tag{35.2-81b}$$

$$\exp(z)\,\frac{\sinh(z)}{z} \;=\; F\left(\begin{matrix}1\\2\end{matrix}\middle|\, 2z\right) \tag{35.2-82a}$$

$$\exp(-iz)\,\frac{\sin(z)}{z} \;=\; F\left(\begin{matrix}1\\2\end{matrix}\middle|\, -2iz\right) \tag{35.2-82b}$$

Further expressions for the sine and cosine are

$$\frac{\sin(a\,z)}{a\,\sin(z)} \;=\; F\left(\begin{matrix}\frac{1+a}{2},\,\frac{1-a}{2}\\ \frac{3}{2}\end{matrix}\middle|\, +\sin(z)^2\right) \tag{35.2-83a}$$

$$\cos(a\,z) \;=\; F\left(\begin{matrix}+\frac{a}{2},\,-\frac{a}{2}\\ \frac{1}{2}\end{matrix}\middle|\, +\sin(z)^2\right) \tag{35.2-83b}$$

$$\frac{\cos(a\,z)}{\cos(z)} \;=\; F\left(\begin{matrix}\frac{1+a}{2},\,\frac{1-a}{2}\\ \frac{1}{2}\end{matrix}\middle|\, +\sin(z)^2\right) \tag{35.2-83c}$$

$$\frac{\sin(a\,z)}{a\,\sin(z)\,\cos(z)} \;=\; \frac{2\,\sin(a\,z)}{a\,\sin(2\,z)} \;=\; F\left(\begin{matrix}1+\frac{a}{2},\,1-\frac{a}{2}\\ \frac{3}{2}\end{matrix}\middle|\, +\sin(z)^2\right) \tag{35.2-83d}$$

Relations for the hyperbolic sine and cosine are obtained by replacing sin $\mapsto$ sinh, cos $\mapsto$ cosh, and negating the sign of the argument of the hypergeometric function. For example, from relation 35.2-83b one obtains

$$\cosh(a\,z) \;=\; F\left(\begin{matrix}+\frac{a}{2},\,-\frac{a}{2}\\ \frac{1}{2}\end{matrix}\middle|\, -\sinh(z)^2\right) \tag{35.2-83e}$$

35.2.5.8 Inverse trigonometric and hyperbolic functions

Series for the inverse tangent and cotangent:

$$\arctan(z) \;=\; -\frac{i}{2}\,\log\frac{1+iz}{1-iz} \;=\; \mathfrak{Im}\log(1+iz) \tag{35.2-84a}$$

$$\;=\; z\,F\left(\begin{matrix}\frac{1}{2},\,1\\ \frac{3}{2}\end{matrix}\middle|\, -z^2\right) \;=\; \sum_{k=0}^{\infty}\frac{(-1)^k\,z^{2k+1}}{2k+1} \tag{35.2-84b}$$

By Pfaff's reflection law (relation 35.2-28b) on can obtain

$$\arctan(z) \;=\; \frac{z}{\sqrt{1+z^2}}\,F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\ \frac{3}{2}\end{matrix}\middle|\, \frac{z^2}{1+z^2}\right) \;=\; \arccos\frac{1}{\sqrt{1+z^2}} \tag{35.2-84c}$$

$$\;=\; \frac{z}{1+z^2}\,F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|\, \frac{z^2}{1+z^2}\right) \quad \text{by 35.2-28b} \tag{35.2-84d}$$

$$\operatorname{arctanh}(z) \;=\; \frac{1}{2}\,\log\frac{1+z}{1-z} \tag{35.2-85a}$$

$$\;=\; z\,F\left(\begin{matrix}\frac{1}{2},\,1\\ \frac{3}{2}\end{matrix}\middle|\, z^2\right) \;=\; \sum_{k=0}^{\infty}\frac{z^{2k+1}}{2k+1} \tag{35.2-85b}$$

$$\operatorname{arccoth}(z) \;=\; \frac{1}{2}\,\log\frac{z+1}{z-1} \;=\; \sum_{k=0}^{\infty}\frac{1}{(2k+1)\,z^{2k+1}} \tag{35.2-85c}$$

$$\log(z) \;=\; 2\,\operatorname{arctanh}\frac{z-1}{z+1} \;=\; 2\,\operatorname{arccoth}\frac{z+1}{z-1} \tag{35.2-85d}$$

$$\begin{aligned}
\operatorname{arccot}(z) &= \arctan\left(\frac{1}{z}\right) = -\frac{i}{2}\log\frac{z+i}{z-i} && (35.2\text{-}86a)\\
&= \frac{1}{z}F\left(\begin{matrix}\frac{1}{2},\,1\\ \frac{3}{2}\end{matrix}\middle|-\frac{1}{z^2}\right) = \sum_{k=0}^{\infty}\frac{(-1)^k}{(2k+1)\,z^{2k+1}} && (35.2\text{-}86b)\\
&= \frac{1}{\sqrt{1+z^2}}F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\ \frac{3}{2}\end{matrix}\middle|\frac{1}{1+z^2}\right) = \arcsin\frac{1}{\sqrt{1+z^2}} && (35.2\text{-}86c)\\
&= \frac{z}{1+z^2}F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|\frac{1}{1+z^2}\right) && (35.2\text{-}86d)
\end{aligned}$$

Applying Clausen's product formula (relation 35.2-46a on page 692) gives

$$\begin{aligned}
[\arctan(z)]^2 &= \frac{z^2}{1+z^2}F\left(\begin{matrix}1,\,1,\,1\\ \frac{3}{2},\,2\end{matrix}\middle|\frac{z^2}{1+z^2}\right) && (35.2\text{-}87a)\\
[\operatorname{arccot}(z)]^2 &= \frac{1}{1+z^2}F\left(\begin{matrix}1,\,1,\,1\\ \frac{3}{2},\,2\end{matrix}\middle|\frac{1}{1+z^2}\right) && (35.2\text{-}87b)
\end{aligned}$$

Series for the inverse sine and cosine are

$$\begin{aligned}
\arcsin(z) &= z\,F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\ \frac{3}{2}\end{matrix}\middle|z^2\right) = \arctan\frac{z}{\sqrt{1-z^2}} && (35.2\text{-}88a)\\
&= z\,F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|\frac{1-\sqrt{1-z^2}}{2}\right) \quad \text{by 35.2-33b} && (35.2\text{-}88b)\\
&= z\sqrt{1-z^2}\,F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|z^2\right) && (35.2\text{-}88c)
\end{aligned}$$

The two latter relations suggest the following argument reduction applicable for the inverse sine (and tangent). Let $G(z) = (1-\sqrt{1-z})/2$, then

$$\begin{aligned}
F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|z\right) &= \frac{1}{\sqrt{1-z}}F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|G(z)\right) && (35.2\text{-}89a)\\
&= \frac{1}{\sqrt{1-z}}\frac{1}{\sqrt{1-G(z)}}F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|G(G(z))\right) = \ldots && (35.2\text{-}89b)\\
F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|z\right) &= \left[\prod_{k=0}^{\infty}1-z_k\right]^{-1/2} \quad \text{where} \quad z_0 = z, \quad z_{k+1} = G(z_k) && (35.2\text{-}89c)
\end{aligned}$$

Clausen's product formula (relation 35.2-46a on page 692) can by applied to obtain

$$[\arcsin(z)]^2 = z^2\,F\left(\begin{matrix}1,\,1,\,1\\ \frac{3}{2},\,2\end{matrix}\middle|z^2\right) \qquad (35.2\text{-}90)$$

$$\arccos(z) = \frac{\pi}{2} - \arcsin(z) = \operatorname{arccot}\frac{z}{\sqrt{1-z^2}} \qquad (35.2\text{-}91)$$

$$\operatorname{arcsinh}(z) = \log(z+\sqrt{1+z^2}) = z\,F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\ \frac{3}{2}\end{matrix}\middle|\,-z^2\right) \tag{35.2-92a}$$

$$= \frac{z}{\sqrt{1+z^2}}\,F\left(\begin{matrix}\frac{1}{2},\,1\\ \frac{3}{2}\end{matrix}\middle|\,\frac{z^2}{1+z^2}\right) \quad \text{by 35.2-28b} \tag{35.2-92b}$$

$$= z\,F\left(\begin{matrix}1,\,1\\ \frac{3}{2}\end{matrix}\middle|\,\frac{1-\sqrt{1+z^2}}{2}\right) \quad \text{by 35.2-33b} \tag{35.2-92c}$$

35.2.6 The function x^x

Boldly setting $a = 1+z$ in $(1+z)^a = F\left(\begin{matrix}-a\end{matrix}\middle|\,-z\right)$ (relation 35.2-62b on page 694) gives

$$(1+z)^{(1+z)} = F\left(\begin{matrix}-1-z\end{matrix}\middle|\,-z\right) = \exp\left[(1+z)\,\log(1+z)\right] \tag{35.2-93a}$$

$$= 1+\frac{(z+1)}{1}\,z\left[1+\frac{(z+0)}{2}\,z\left[1+\frac{(z-1)}{3}\,z\left[1+\frac{(z-2)}{4}\,z\left[1+\ldots\ \right]\right]\right]\right] \tag{35.2-93b}$$

$$= 1+z+z^2+\frac{1}{2}\,z^3+\frac{1}{3}\,z^4+\frac{1}{12}\,z^5+\frac{3}{40}\,z^6-\frac{1}{120}\,z^7+\frac{59}{2520}\,z^8-\frac{71}{5040}\,z^9\pm\ldots \tag{35.2-93c}$$

This somewhat surprising expression allows the computation of x^x without computing exp() or log(). The series converges for real $z>0$ so we can compute x^x (where $x=1+z$) for real $x>+1$ as

$$x^x = F\left(\begin{matrix}-x\end{matrix}\middle|\,-x+1\right) \tag{35.2-94a}$$

$$= 1+\frac{x-0}{1}\,(x-1)\left[1+\frac{x-1}{2}\,(x-1)\left[1+\frac{x-2}{3}\,(x-1)\left[1+\ldots\ \right]\right]\right] \tag{35.2-94b}$$

We denote the series obtained by truncating after the n-th term of the hypergeometric function by $g_n(x)$. For example, with $n=2$ and $n=4$ we obtain:

$$g_2(x) = \frac{1}{2}\,x^4-\frac{3}{2}\,x^3+\frac{5}{2}\,x^2-\frac{3}{2}\,x+1 \tag{35.2-95a}$$

$$= \frac{1}{2}\,z^4+\frac{1}{2}\,z^3+z^2+z+1 \tag{35.2-95b}$$

$$g_4(x) = \frac{1}{24}\,x^8-\frac{5}{12}\,x^7+\frac{15}{8}\,x^6-\frac{19}{4}\,x^5+\frac{61}{8}\,x^4-\frac{31}{4}\,x^3+\frac{131}{24}\,x^2-\frac{25}{12}\,x+1 \tag{35.2-95c}$$

$$= \frac{1}{24}\,z^8-\frac{1}{12}\,z^7+\frac{1}{8}\,z^6+\frac{1}{12}\,z^5+\frac{1}{3}\,z^4+\frac{1}{2}\,z^3+z^2+z+1 \tag{35.2-95d}$$

We have $g_n(n)=n^n$ and further $g_n(k)=k^k$ for all integer $k\le n$. Thus we have just invented a curious way to find polynomials of degree $2n$ that interpolate k^k for $0\le k\le n$ (setting $0^0:=1$ for our purposes).

The polynomials actually give acceptable estimates for x^x also for non-integer x, especially for x near 1. The (unique) degree-n polynomials $i_n(x)$ that are obtained by interpolating the values k^k have much bigger coefficients and give values far away from x^x for non-integer arguments x.

For $0<x<n$ the interpolating polynomials $i_n(x)$ give an estimate that is consistently worse than $g_n(x)$ for non-integer values of x. The same is true even for the polynomials $i_{2n}(x)$ that interpolate k^k for $0\le k\le 2n$ (so that $\deg(i_{2n})=\deg(g_n)=2n$). In fact, the $i_{2n}(x)$ approximate consistently worse than $i_n(x)$ for non-integer x.

Finally, the Padé approximants $p_{[n,n]}(x)$ for $g_n(x)$ give estimates that are worse than with both $i_n(x)$ or $g_n(x)$. Further, $g_n(x)\neq x^x$ even for integer x and the $p_{[n,n]}(x)$ have a pole on the real axis near $x=1$. That is, we found a surprisingly good and compact polynomial approximation for the function x^x.

The sequence of the n-th derivatives of $(1+z)^{(1+z)}$ at $z=0$ is entry A005727 of [240]:

```
? Vec(serlaplace(exp((1+z)*log(1+z))))
  [1, 1, 2, 3, 8, 10, 54, -42, 944, -5112, 47160, -419760, 4297512, ... ]
```

Many other expressions for the function x^x can be given, we note just one: set $s = 2+z$ in relation 35.2-64 on page 694 to obtain

$$(1+z)^{(1+z)} = \frac{1}{1+2\,z}\,F\left(\begin{matrix}-z-2,\ z+3\\ z+2\end{matrix}\,\middle|\,-z\right) \tag{35.2-96}$$

35.2.7 Elliptic K and E

In order to avoid the factor $\frac{\pi}{2}$ we let $\tilde{K} := \frac{2\,K}{\pi}$, $\tilde{E} := \frac{2\,E}{\pi}$, then

$$\tilde{K}(k) = F\left(\begin{matrix}\frac{1}{2},\ \frac{1}{2}\\ 1\end{matrix}\,\middle|\,k^2\right) \tag{35.2-97a}$$

$$\tilde{E}(k) = F\left(\begin{matrix}-\frac{1}{2},\ \frac{1}{2}\\ 1\end{matrix}\,\middle|\,k^2\right) \tag{35.2-97b}$$

See section 30.2 on page 595 for explicit expansions. We further set

$$\tilde{N}(k) = F\left(\begin{matrix}-\frac{1}{2},\ -\frac{1}{2}\\ 1\end{matrix}\,\middle|\,k^2\right) \tag{35.2-97c}$$

$$= 1+\frac{1}{4}\,k^2+\frac{1}{64}\,k^4+\frac{1}{256}\,k^6+\frac{25}{16384}\,k^8+\frac{49}{65536}\,k^{10}+\frac{441}{1048576}\,k^{10}+\ldots \tag{35.2-97d}$$

A special value is $\tilde{N}(1) = 4/\pi$. Most of the following relations can be written in several ways by one of the identities

$$k' = \sqrt{1-k^2} \tag{35.2-98a}$$

$$-\frac{k^2}{1-k^2} = -\left(\frac{k}{k'}\right)^2 = -\frac{1-k'^2}{k'^2} \tag{35.2-98b}$$

$$-\left(\frac{k}{1-\sqrt{1-k^2}}\right)^2 = -\frac{1+k'}{1-k'} = -\frac{1-k'^2}{(1-k')^2} = -\frac{(1+k')^2}{1-k'^2} \tag{35.2-98c}$$

$$\frac{2}{1+k'} = 1+\frac{1-k'}{1+k'} \tag{35.2-98d}$$

$$\frac{4k}{(1+k)^2} = 1-\left(\frac{1-k}{1+k}\right)^2 \tag{35.2-98e}$$

$$\frac{-4k^2}{(1-k^2)^2} = -\left(\frac{2k}{k'^2}\right)^2 = 1-\left(\frac{1+k^2}{1-k^2}\right)^2 \tag{35.2-98f}$$

35.2.7.1 Relations for K

$$\tilde{K}(k) = F\left(\begin{matrix}\frac{1}{2},\ \frac{1}{2}\\ 1\end{matrix}\,\middle|\,k^2\right) \tag{35.2-99}$$

$$= \frac{1}{k'}\,F\left(\begin{matrix}\frac{1}{2},\ \frac{1}{2}\\ 1\end{matrix}\,\middle|\,-\left(\frac{k}{k'}\right)^2\right) \qquad \text{by 35.2-28b} \tag{35.2-100}$$

$$\tilde{K}(k) = \frac{1}{1-k'}\,F\left(\begin{matrix}\frac{1}{2},\ \frac{1}{2}\\ 1\end{matrix}\,\middle|\,-\frac{1+k'}{1-k'}\right) \tag{35.2-101}$$

From the product form (relation 30.2-8 on page 596) we obtain

$$\tilde{K}(k) \quad = \quad \frac{2}{1+k'}\, F\left(\begin{matrix}\frac{1}{2},\,\frac{1}{2}\\ 1\end{matrix}\,\middle|\, \left(\frac{1-k'}{1+k'}\right)^2\right) \tag{35.2-102a}$$

The relation can be written as

$$\tilde{K}(k) \quad = \quad (1+z(k))\;\tilde{K}\left(z(k)\right) \quad \text{where} \quad z(k) := \frac{1-k'}{1+k'} \tag{35.2-102b}$$

Relation 35.2-42f on page 691 with $a = \frac{1}{2}$ gives

$$\tilde{K}(k) \quad = \quad \frac{1}{1+k}\,\tilde{K}\left(\frac{2\sqrt{k}}{1+k}\right) \tag{35.2-103}$$

$$\tilde{K}(k) \quad = \quad \frac{1}{\sqrt{1-k^2}}\, F\left(\begin{matrix}\frac{1}{4},\,\frac{1}{4}\\ 1\end{matrix}\,\middle|\, \frac{-4k^2}{(1-k^2)^2}\right) \;=\; \frac{1}{k'}\, F\left(\begin{matrix}\frac{1}{4},\,\frac{1}{4}\\ 1\end{matrix}\,\middle|\, 1-\left(\frac{1+k^2}{1-k^2}\right)^2\right) \tag{35.2-104a}$$

Euler's transform on 35.2-104a gives:

$$\tilde{K}(k) \quad = \quad \frac{1}{k'}\,\frac{1+k^2}{1-k^2}\, F\left(\begin{matrix}\frac{3}{4},\,\frac{3}{4}\\ 1\end{matrix}\,\middle|\, -\left(\frac{2k}{k'^2}\right)^2\right) \tag{35.2-105}$$

$$\tilde{K}(k) \quad = \quad F\left(\begin{matrix}\frac{1}{4},\,\frac{1}{4}\\ 1\end{matrix}\,\middle|\, (2\,k\,k')^2\right) \qquad \text{by 35.2-33a} \tag{35.2-106}$$

35.2.7.2 Relations for E

$$\tilde{E}(k) \quad = \quad F\left(\begin{matrix}-\frac{1}{2},\,\frac{1}{2}\\ 1\end{matrix}\,\middle|\, k^2\right) \tag{35.2-107}$$

$$\tilde{E}(k) \quad = \quad k'\, F\left(\begin{matrix}-\frac{1}{2},\,\frac{1}{2}\\ 1\end{matrix}\,\middle|\, -\left(\frac{k}{k'}\right)^2\right) \qquad \text{by 35.2-28b} \tag{35.2-108}$$

The following relation resembles relation 35.2-102a:

$$\tilde{E}(k) \quad = \quad \frac{\left(1+\sqrt{1-k^2}\right)}{2}\, F\left(\begin{matrix}-\frac{1}{2},\,-\frac{1}{2}\\ 1\end{matrix}\,\middle|\, \left(\frac{1-\sqrt{1-k^2}}{1+\sqrt{1-k^2}}\right)^2\right) \tag{35.2-109a}$$

$$= \quad \frac{1+k'}{2}\, F\left(\begin{matrix}-\frac{1}{2},\,-\frac{1}{2}\\ 1\end{matrix}\,\middle|\, \left(\frac{1-k'}{1+k'}\right)^2\right) \tag{35.2-109b}$$

$$= \quad (1+z(k))^{-1}\;\tilde{N}\left(z(k)\right) \quad \text{where} \quad z(k) := \frac{1-k'}{1+k'} \tag{35.2-109c}$$

35.2.7.3 Relations for N

$$\tilde{N}(k) \quad = \quad \sqrt{1-k^2}\, F\left(\begin{matrix}-\frac{1}{4},\,\frac{3}{4}\\ 1\end{matrix}\,\middle|\, \frac{-4\,k^2}{(1-k^2)^2}\right) \qquad \text{by 35.2-42e} \tag{35.2-110a}$$

$$= \quad k'\, F\left(\begin{matrix}-\frac{1}{4},\,\frac{3}{4}\\ 1\end{matrix}\,\middle|\, -\left(\frac{2k}{k'^2}\right)^2\right) \tag{35.2-110b}$$

$$= \quad \sqrt{1+k^2}\, F\left(\begin{matrix}-\frac{1}{4},\,\frac{1}{4}\\ 1\end{matrix}\,\middle|\, \left(\frac{2k}{k^2+1}\right)^2\right) \qquad \text{by 35.2-28b} \tag{35.2-110c}$$

Relation 35.2-25 on page 689 with $a = b = -1/2$ and $c = 1$ gives

$$F\left(\begin{matrix}\frac{1}{2}, -\frac{1}{2} \\ 1\end{matrix}\middle| z\right) = \frac{1}{2} F\left(\begin{matrix}\frac{3}{2}, -\frac{1}{2} \\ 1\end{matrix}\middle| z\right) + \frac{1}{2} F\left(\begin{matrix}\frac{1}{2}, \frac{1}{2} \\ 1\end{matrix}\middle| z\right) \tag{35.2-111}$$

Applying 35.2-28c on page 689 to the second function ($F\left(\begin{matrix}3/2, -1/2 \\ 1\end{matrix}\middle| z\right) = (1-z)\, F\left(\begin{matrix}-1/2, -1/2 \\ 1\end{matrix}\middle| \frac{-z}{1-z}\right)$) and rearranging gives

$$2\,\tilde{E}(k) - \tilde{K}(k) = k'\,\tilde{N}\left(i\frac{k}{k'}\right) \tag{35.2-112}$$

Applying the transformation 35.2-42f on page 691 on the defining relation gives the key to fast computation of the function $\tilde{N}(k)$:

$$\tilde{N}(k) = (1+k)\,\tilde{E}\left(\frac{2\sqrt{k}}{1+k}\right) \tag{35.2-113}$$

The relation

$$2\,\tilde{E}(k) - k'^2\,\tilde{K}(k) = \tilde{N}(k) \tag{35.2-114}$$

can be used to rewrite Legendre's relation (equation 30.2-17b on page 597) as either

$$\frac{4}{\pi} = \tilde{N}\,\tilde{K}' + \tilde{K}\,\tilde{N}' - \tilde{K}\,\tilde{K}' \tag{35.2-115}$$

or, by setting $N := \pi/2\,\tilde{N}$,

$$\pi = N\,K' + K\,N' - K\,K' \tag{35.2-116}$$

35.3 Continued fractions

A *continued fraction* is an expression of the form:

$$K(a,b) = a_0 + \cfrac{b_1}{a_1 + \cfrac{b_2}{a_2 + \cfrac{b_3}{a_3 + \cfrac{b_4}{a_4 + \cdots}}}} \tag{35.3-1}$$

Continued fractions are sometimes expressed in the following form:

$$K(a,b) = a_0 + \frac{b_1}{a_1+}\,\frac{b_2}{a_2+}\,\frac{b_3}{a_3+}\,\frac{b_4}{a_4+}\,\cdots \tag{35.3-2}$$

The a_k and b_k are called the k-th *partial numerators* and *denominators*.

For $k > 0$ let P_k/Q_k be the value of the above fraction if b_{k+1} is set to zero (that is, the continued fraction terminates at index k). The ratio is called the k-th *convergent* of the continued fraction:

$$\frac{P_k}{Q_k} = a_0 + \frac{b_1}{a_1+}\,\frac{b_2}{a_2+}\,\frac{b_3}{a_3+}\,\frac{b_4}{a_4+}\,\cdots\,\frac{b_{k-1}}{a_{k-1}+}\,\frac{b_k}{a_k} \tag{35.3-3}$$

We note that multiplication of a_i, b_i, b_{i+1} by some nonzero value does not change the value of the continued fraction. Thereby

$$a_0 + \frac{b_1}{a_1+}\,\frac{b_2}{a_2+}\,\frac{b_3}{a_3+}\,\frac{b_4}{a_4+}\,\cdots = a_0 + \frac{c_1\,b_1}{c_1\,a_1+}\,\frac{c_2\,c_1\,b_2}{c_2\,a_2+}\,\frac{c_3\,c_2\,b_3}{c_3\,a_3+}\,\frac{c_4\,c_3\,b_4}{c_4\,a_4+}\,\cdots \tag{35.3-4}$$

where all c_i are arbitrary nonzero constants.

35.3.1 Simple continued fractions

Continued fractions where all b_k are equal to one (and all the a_k are positive) are called *simple continued fractions*. Rational numbers have terminating continued fractions. Note that the expression of a rational number as simple continued fraction is not unique:

$$[a_0,\, \ldots,\, a_{n-1},\, a_n] \;=\; [a_0,\, \ldots,\, a_{n-1},\, a_n - 1,\, 1] \quad \text{if } a_n > 1 \tag{35.3-5a}$$

$$[a_0,\, \ldots,\, a_{n-1},\, 1] \;=\; [a_0,\, \ldots,\, a_{n-1} + 1] \qquad \text{if} \quad a_n = 1 \tag{35.3-5b}$$

Solutions of the quadratic equation $\alpha\, x^2 + \beta\, x + \gamma = 0$ that are not rational ($\Delta := \beta^2 - 4\,\alpha\,\gamma$ not a square) have simple continued fractions that are eventually periodic. For example:

```
? contfrac(sqrt(5))
  [2, 4, 4, 4, 4, 4, ...]
? contfrac(2+sqrt(3))
  [3, 1,2, 1,2, 1,2, ...]
? contfrac(sqrt(19))
  [4, 2,1,3,1,2,8, 2,1,3,1,2,8, 2,1,3,1,2,8, ...]
```

For the k-th convergent P_k/Q_k (in lowest terms) of the simple continued fraction expansion of some number x then the convergent is a *best approximation* in the following sense: if p/q is any better rational approximation to x (that is, $\left|\frac{p}{q} - x\right| < \left|\frac{P_k}{Q_k} - x\right|$), then one must have $q > Q_k$.

For the simple continued fraction of x one has

$$\left| x - \frac{P_n}{Q_n} \right| \;\leq\; \frac{1}{Q_n\, Q_{n-1}} \;<\; \frac{1}{Q_n^2} \tag{35.3-6}$$

and equality can only occur with terminating continued fractions.

35.3.1.1 Computing the simple continued fraction of a number

Given a numerical quantity one can compute the sequence a_k of its simple continued fraction by the following algorithm:

```
1   procedure number_to_scf(x, n, a[0..n-1])
2   {
3       for k:=0 to n-1
4       {
5           xi := floor(x)
6           a[k] := xi
7           x := 1 / (x-xi)
8       }
9   }
```

Here `n` is the number of requested terms a_k. Obviously some check has to be inserted in order to avoid possible division be zero (and indicate a terminating continued fraction as will occur for rational x). If in the process one keeps track of the exact rational values of the simple continued fraction convergents (using the recursion relations) then the algorithm can be used to produce a 'best possible' rational approximation where the denominator does not exceed a certain specified size. Alternatively one can stop as soon as the convergent is within some specified bound.

35.3.1.2 Continued fractions of polynomial roots

Let $r > 1$ be the only real positive root of a polynomial $F(x)$ with integer coefficients and positive leading coefficient. Then the (simple) continued fraction $[a_0, a_1, \ldots, a_n]$ of r can be computed as follows (taken from [174, p.261]):

1. Set $k = 0$, $F_0(x) = F(x)$, and $d = \deg(F)$.
2. Find the (unique) real positive root r_k of $F_k(x)$, set $a_k = \lfloor r \rfloor$. If $k = n$ then stop.

```
r = RootOf(z^3 - 2) == 1.2599210...
contfrac(r) == [1, 3, 1, 5, 1, 1, 4, 1, 1, 8, 1, 14, 1, 10, 2, ...]

f=z^3 - 2    r=1.25992104989487  ==> 1
f=z^3 - 3*z^2 - 3*z - 1    r=3.84732210186307  ==> 3
f=10*z^3 - 6*z^2 - 6*z - 1    r=1.18018873554841  ==> 1
f=3*z^3 - 12*z^2 - 24*z - 10    r=5.54973648578239  ==> 5
f=55*z^3 - 81*z^2 - 33*z - 3    r=1.81905335713127  ==> 1
f=62*z^3 + 30*z^2 - 84*z - 55    r=1.22092167902528  ==> 1
f=47*z^3 - 162*z^2 - 216*z - 62    r=4.52649103705930  ==> 4
f=510*z^3 - 744*z^2 - 402*z - 47    r=1.89936756679748  ==> 1
f=683*z^3 + 360*z^2 - 786*z - 510    r=1.11189244188653  ==> 1
f=253*z^3 - 1983*z^2 - 2409*z - 683    r=8.93715413784671  ==> 8
f=17331*z^3 - 14439*z^2 - 4089*z - 253    r=1.06706032616757  ==> 1
f=1450*z^3 - 19026*z^2 - 37554*z - 17331    r=14.9119465584038  ==> 14
```

Figure 35.3-A: Computation of the continued fraction of the positive real root of the polynomial $z^3 - 2$.

```
r = RootOf(z^2 - 29) == 5.38516480...
contfrac(r) ==[5,   2, 1, 1, 2, 10,   2, 1, 1, 2, 10,  ...]

f=z^2 - 29    r=5.38516480713450  ==> 5
f=4*z^2 - 10*z - 1    r=2.59629120178363  ==> 2
f=5*z^2 - 6*z - 4    r=1.67703296142690  ==> 1
f=5*z^2 - 4*z - 5    r=1.47703296142690  ==> 1
f=4*z^2 - 6*z - 5    r=2.09629120178363  ==> 2
f=z^2 - 10*z - 4    r=10.3851648071345  ==> 10

f=4*z^2 - 10*z - 1    r=2.59629120178363  ==> 2  <--= restart period
f=5*z^2 - 6*z - 4    r=1.67703296142690  ==> 1
```

Figure 35.3-B: Computation of the continued fraction of the positive real root of the polynomial $z^2 - 29$.

3. Set $G(x) = F_k(x + a_k)$, set $F_{k+1} = -G^*(x) = -x^d\, G(1/x)$.

4. Set $k = k + 1$ and goto step 2.

A simple demonstration is

```
1   f = z^3 - 2
2   ff(y)=subst(f, z, y)  \\ for solve() function
3
4   { for (k=1, 12,
5       print1("  f=", f);
6       r = solve(x=0.9, 1e9, ff(x));  \\ lazy implementation
7       print1("    r=", r);
8       ak = floor( r );
9       print1("  ==> ", ak);
10      g = subst(f, z, z+ak);  \\ shifted polynomial
11      f = -polrecip( g );     \\ negated reciprocal of g
12      print();
13   ); }
```

The output with $F(x) = x^3 - 2$ is shown in figure 35.3-A. With quadratic equations one obtains periodic continued fractions, figure 35.3-B shows the computation for $F(x) = x^2 - 29$. For a comparison of methods for the computation of continued fractions for algebraic numbers see [68].

35.3.2 Computation of the convergents (evaluation)

The computation of the sequence of convergents uses the recurrence

$$P_k = a_k\, P_{k-1} + b_k\, P_{k-2} \tag{35.3-7a}$$

$$Q_k = a_k\, Q_{k-1} + b_k\, Q_{k-2} \tag{35.3-7b}$$

Set $\frac{P_{-1}}{Q_{-1}} := \frac{1}{0}$ and $\frac{P_0}{Q_0} := \frac{a_0}{1}$ to initialize. The following is a procedure that computes the sequences of values P_k and Q_k for $k = -1 \ldots n$ for a given continued fraction:

```
1   procedure ratios_from_cf(a[0..n], b[0..n], n, P[-1..n], Q[-1..n])
2   {
3       P[-1] := 1
4       Q[-1] := 0
5
6       P[0] := a[0]
7       Q[0] := 1
8
9       for k:=1 to n
10      {
11          P[k] := a[k] * P[k-1] + b[k] * P[k-2]
12          Q[k] := a[k] * Q[k-1] + b[k] * Q[k-2]
13      }
14  }
```

If only the last ratio is of interest, the 'backward' variant of the algorithm can be used. The version that computes the numerical value `x` from the first `n` terms of a simple continued fraction is:

```
1   function ratio_from_cf(a[0..n-1], n)
2   {
3       x := a[n-1]
4       for k:=n-2 to 0 step -1
5       {
6           x := 1/x + a[k]
7       }
8
9       return x
10  }
```

Using rational arithmetic and a general (non-simple) continued fraction, the algorithm becomes:

```
1   function ratio_from_cf(a[0..n-1], b[0..n-1], n)
2   {
3       P := a[n-1]
4       Q := b[n-1]
5       for k:=n-2 to 0 step -1
6       {
7           {P, Q} := {a[k]*P+b[k]*Q, P}  // x := b[k] / x + a[k]
8       }
9
10      return P/Q
11  }
```

Implementation

Converting a number to a simple continued fraction can be done with pari/gp's builtin function `contfrac()`. The final convergent can be computed with `contfracpnqn()`:

```
? default(realprecision,23)
    realprecision = 28 significant digits (23 digits displayed)
? Pi
    3.1415926535897932384626
? cf=contfrac(Pi)
    [3, 7, 15, 1, 292, 1, 1, 1, 2, 1, 3, 1, 14, 2, 1, 1, 2, 2, 2, 2, 1, 84, 2, 1, 1, 15, 3]
? ?contfracpnqn
    contfracpnqn(x): [p_n,p_{n-1}; q_n,q_{n-1}] corresponding to the continued fraction x.
? m=contfracpnqn(cf)
    [428224593349304 139755218526789]
    [136308121570117 44485467702853]
? 1.0*m[1,1]/m[2,1]
    3.1415926535897932384626
```

The number of terms of the continued fraction depends on the precision used, with greater precision more terms can be computed. The computation of the m-th convergent of a continued fraction given as two vectors `a[]` and `b[]` can be implemented as (backward variant):

```
1   cfab2r(a,b, m=-2)=
2   {
3       local(n, r);
4       n = length(a);
5       if ( m>-2, m = min(n, m) ); \\ default: m=n
6       if ( m>=n,  m=n-1 );
7       if ( m<0, return( 0 ) ); \\ infinity
8       r = 0;
9       m += 1;
10      forstep (k=m, 2, -1,  r = b[k]/(a[k]+r); );
```

```
11      r += a[1];  \\ b[1] unused
12      return( r );
13  }
```

Alternatively, one can use the recursion relations 35.3-7a and 35.3-7b. We do not store all pairs P_n, Q_n but only return the final pair P_m, Q_m:

```
1   cfab2pq(a,b,m=-2)=
2   {
3       local(n, p, p1, p2, q, q1, q2, i);
4       n = length(a)-1;
5       if ( m>-2, m = min(n, m) ); \\ default: m=n
6       if ( m<0, return( [1, 0] ) ); \\ infinity
7       p1 = 1;
8       q1 = 0;
9       p = a[1];
10      q = 1;  \\ b[1] unused
11      for (k=1, m,
12          i = k+1;
13          p2 = p1;  p1 = p;
14          q2 = q1;  q1 = q;
15          p = a[i]*p1 + b[i]*p2;
16          q = a[i]*q1 + b[i]*q2;
17      );
18      return( [p,q] );
19  }
```

```
default(realprecision, 55);  \\ use enough precision
default(format, "g.11"); \\ print with moderate precision
default(echo, 0);
\r contfrac.inc.gp  \\ functions cfab2pq and cfab2r

x=4.0/Pi
n=15
/* set up the continued fraction: */
a=vector(n, j, 2);  a[1]=1;
b=vector(n, j, (2*j-3)^2);

/* print convergents and their error: */
{ for(k=0, n-1,
    t=cfab2pq(a,b, k);
    p=t[1]; q=t[2];
    print1(k, ":  ",p, " / ", q);
    print1("\n  d=", x-p/q);
    print();
  ); }

quit; /* ------ end of script -------- */
/* ------ start output: ------ */
15  /* =n */
[1, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2]  /* =a */
[1, 1, 9, 25, 49, 81, 121, 169, 225, 289, 361, 441, 529, 625, 729]  /* =b */
1.2732395447 /* = 4/Pi */

0:  1 / 1
  d=0.27323954473
1:  3 / 2
  d=-0.22676045526
2:  15 / 13
  d=0.11939339088
3:  105 / 76          /* =p3/q3 */
  d=-0.10833940263    /* =p3/q3-4/Pi */
4:  945 / 789
  d=0.075520913556
5:  10395 / 7734
  d=-0.070825622061
[--snip--]
13:  213458046676875 / 163842638377950
  d=-0.029583998575
14:  6190283353629375 / 4964894559637425
  d=0.026428906710
```

Figure 35.3-C: A pari/gp script demonstrating the function `cfab2pq()` that computes the convergents of a continued fraction (top) and its output (bottom, comments added). Here convergence is rather slow.

We use our routines to compute the convergents of the continued fraction for $4/\pi$ given 1658 by Brouncker:

$$\frac{4}{\pi} \;=\; 1+\cfrac{1^2}{2+\cfrac{3^2}{2+\cfrac{5^2}{2+\cfrac{7^2}{2+\cdots}}}} \;=\; a_0+\cfrac{b_1}{a_1+\cfrac{b_2}{a_2+\cfrac{b_3}{a_3+\cfrac{b_4}{a_4+\cdots}}}} \tag{35.3-8}$$

Figure 35.3-C shows how to set up the vectors containing the a_k and b_k and check the convergents.

Fast evaluation as matrix product

For the evaluation of a continued fraction with a large number of terms rewrite relations 35.3-7a and 35.3-7b as

$$\begin{bmatrix} P_k & Q_k \\ P_{k-1} & Q_{k-1} \end{bmatrix} \;=\; \prod_{j=0}^{k} \begin{bmatrix} a_j & 1 \\ b_j & 0 \end{bmatrix} \tag{35.3-9}$$

Use the binary splitting algorithm (section 32.1 on page 631) for the computation of the matrix product $M(k)$.

35.3.3 Miscellaneous relations for continued fractions

35.3.3.1 Determinantal expressions

The *determinant formula* for the numerators and denominators of successive convergents is

$$\det \begin{bmatrix} P_k & Q_k \\ P_{k-1} & Q_{k-1} \end{bmatrix} \;=\; P_k\,Q_{k-1} - P_{k-1}\,Q_k \;=\; (-1)^{k-1} \prod_{j=1}^{k} b_j \tag{35.3-10}$$

The relation is obtained by taking determinants on both sides of equation 35.3-9. The relation can also be written as

$$\frac{P_k}{Q_k} - \frac{P_{k-1}}{Q_{k-1}} \;=\; \frac{(-1)^{k-1} \prod_{j=1}^{k} b_j}{Q_{k-1}\,Q_k} \tag{35.3-11}$$

For simple continued fractions we have $b_j = 1$ so the product in the numerator equals one. Further, by inserting $P_{k-1} = (P_k - b_k\,P_{k-2})/a_k$ (relation 35.3-7a) and the equivalent expression for Q_{k-1} into relation 35.3-10, one obtains

$$\det \begin{bmatrix} P_k & Q_k \\ P_{k-2} & Q_{k-2} \end{bmatrix} \;=\; P_k\,Q_{k-2} - P_{k-2}\,Q_k \;=\; (-1)^{k}\,a_k \prod_{j=1}^{k-1} b_j \tag{35.3-12}$$

Equivalently,

$$\frac{P_k}{Q_k} - \frac{P_{k-2}}{Q_{k-2}} \;=\; \frac{(-1)^{k}\,a_k \prod_{j=1}^{k-1} b_j}{Q_{k-2}\,Q_k} \tag{35.3-13}$$

This relation tells us (provided all a_j and b_j are positive) that the sequence of even convergents is increasing and the sequence of odd convergents is decreasing. As both converge to a common limit we have $P_o/Q_o \geq P_e/Q_e$ for all even e and odd o. Equality can occur only for terminating continued fractions.

35.3.3.2 Subsequences of convergents

Sometimes the terms a_k, b_k of the continued fraction are given in the form "$a_k = u(k)$ *if* k *even,* $a_k = v(k)$ *else*" (and b_k equivalently). Then one might want to compute the $x = K(a, b)$ in a 'stride two' manner as

$$P_k = A_k\, P_{k-2} + B_k\, P_{k-4} \tag{35.3-14a}$$
$$Q_k = A_k\, Q_{k-2} + B_k\, Q_{k-4} \tag{35.3-14b}$$

in order to regularize the involved expressions. We write the recurrence relation three times

$$P_k = a_k\, P_{k-1} + b_k\, P_{k-2} \tag{35.3-15a}$$
$$P_{k-1} = a_{k-1}\, P_{k-2} + b_{k-1}\, P_{k-3} \tag{35.3-15b}$$
$$P_{k-2} = a_{k-2}\, P_{k-3} + b_{k-2}\, P_{k-4} \tag{35.3-15c}$$

and eliminate the terms P_{k-1} and P_{k-3}. This gives

$$A_k = \frac{a_k\, b_{k-1} + b_k\, a_{k-2} + a_k\, a_{k-1}\, a_{k-2}}{a_{k-2}} = \frac{a_k\, b_{k-1}}{a_{k-2}} + b_k + a_k\, a_{k-1} \tag{35.3-16a}$$
$$B_k = \frac{-a_k\, b_{k-1}\, b_{k-2}}{a_{k-2}} \tag{35.3-16b}$$

The stride three version

$$P_k = A_k\, P_{k-3} + B_k\, P_{k-6} \tag{35.3-17a}$$
$$Q_k = A_k\, Q_{k-3} + B_k\, Q_{k-6} \tag{35.3-17b}$$

leads to the expressions (writing a_n for a_{k-n} to reduce line width):

$$A_k = \frac{a_0b_1b_3 + b_0a_2b_3 + b_0b_2a_4 + a_0a_1a_2b_3 + a_0a_1b_2a_4 + a_0b_1a_3a_4 + b_0a_2a_3a_4 + a_0a_1a_2a_3a_4}{b_3 + a_3a_4} \tag{35.3-18a}$$
$$B_k = \frac{b_0b_2b_3b_4 + a_0a_1b_2b_3b_4}{b_3 + a_3a_4} \tag{35.3-18b}$$

When setting $a_k = \alpha$, $b_k = \beta$ the expressions for A_k and B_k simplify to the coefficients in relations 34.1-13c on page 663 (stride two) and 34.1-13d (stride three) for recurrences.

35.3.3.3 Relation to alternating series

Using relation 35.3-11 it is possible to rewrite a continued fraction $x = K(a, b)$ with positive a_k, b_k as an alternating series

$$x = a_0 + \sum_{k=1}^{\infty} (-1)^{k+1}\, s_k \tag{35.3-19}$$
$$= a_0 + \frac{b_1}{Q_0\, Q_1} - \frac{b_1\, b_2}{Q_1\, Q_2} + \frac{b_1\, b_2\, b_3}{Q_2\, Q_3} \pm \ldots + (-1)^{k+1}\, \frac{\prod_{i=1}^{k} b_i}{Q_k\, Q_{k+1}} \pm \ldots$$

Thereby the algorithm for the accelerated summation of alternating series from section 32.3 can be applied to compute x.

35.3.3.4 Continued fractions from infinite products

A continued fraction for the product

$$P := \prod_{k=0}^{\infty} (1 + Y_k) \tag{35.3-20a}$$

in terms of $a = [a_0, a_1, \ldots]$ and $b = [b_0, b_1, \ldots]$ is

$$a = [1, \quad 1, \quad +Y_1 + (1+Y_1)\cdot Y_0, \quad +Y_2 + (1+Y_2)\cdot Y_1, \quad +Y_3 + (1+Y_3)\cdot Y_2, \ \ldots\,] \quad \text{(35.3-20b)}$$

$$b = [1, \quad +Y_0, \quad -1\cdot Y_1\cdot(1+Y_0), \quad -Y_0\cdot Y_2\cdot(1+Y_1), \quad -Y_1\cdot Y_3\cdot(1+Y_2), \ \ldots\,] \quad \text{(35.3-20c)}$$

```
Y(k) = eval(Str("Y" k))  \\ return symbol Yk
yprod(n)= if (n<=0, 1, prod(j=0, n-1, (1+Y(j))))
n=3
pr = yprod(n)
    ((Y2 + 1)*Y1 + (Y2 + 1))*Y0 + ((Y2 + 1)*Y1 + (Y2 + 1))
yv = vector(n, j, Y(j-1) )
    [Y0, Y1, Y2]
t = cfprod(yv);
a=t[1]
    [1, 1, (Y1 + 1)*Y0 + Y1, (Y2 + 1)*Y1 + Y2]
b=t[2]
    [1, Y0, -Y1*Y0 - Y1, (-Y2*Y1 - Y2)*Y0]
{ for(k=0, n,
    t=cfab2pq(a,b, k);
    p=t[1]; q=t[2];
    print1(k, ":  (",p, ") / (", q,")");
    yp = yprod(k);
    print1("\n   == ", simplify(p/q));
    print();
  ); }
    0:  (1) / (1)
       == 1
    1:  (Y0 + 1) / (1)  \\  (p1) / (q1)
       == Y0 + 1        \\  == yprod(1)
    2:  ((Y1 + 1)*Y0^2 + (Y1 + 1)*Y0) / (Y0)  \\  (p2) / (q2)
       == (Y1 + 1)*Y0 + (Y1 + 1)              \\ == yprod(2) == (1+Y0)*(1+Y1)
    3:  (((Y2 + 1)*Y1^2 + (Y2 + 1)*Y1)*Y0^2 + ((Y2 + 1)*Y1^2 + (Y2 + 1)*Y1)*Y0) / (Y1*Y0)
       == ((Y2 + 1)*Y1 + (Y2 + 1))*Y0 + ((Y2 + 1)*Y1 + (Y2 + 1))
```

Figure 35.3-D: Verification of relations 35.3-20b and 35.3-20c using pari/gp.

For a given a vector $y = [Y_0, Y_1, \ldots]$ the computation of individual values a_k and b_k can be implemented as:

```
1   cfproda(yv, n)=
2   {
3       local( y2, y3, d );
4       if ( n<=1, return(1) );
5       y3 = yv[n];
6       y2 = yv[n-1];
7       return( (y2+y3*(1+y2)) );
8
9   }
```

```
1   cfprodb(yv, n)=
2   {
3       local( y1, y2, y3 );
4       if (0==n, return(1) );  \\ unused
5       if (1==n, return(+yv[1+0]) );
6       y3 = yv[n];
7       y2 = yv[n-1];
8       y1 = if ( n==2, 1, yv[n-2] );
9       return( -y1*y3*(1+y2) );
10
11  }
```

The routine `cfprod()` generates the vectors a and b with $n+1$ terms where n is the length of y:

```
1   cfprod(yv)=
2   {
3       local(n, a, b);
4       n = length(yv);
5       n += 1;  \\ n+1 terms in continued fraction
6       a = vector(n);
7       b = vector(n);
```

```
8       for (k=0, n-1,
9           a[k+1] = cfproda(yv, k);
10          b[k+1] = cfprodb(yv, k);
11      );
12      return( [a, b] );
13  }
```

Relations 35.3-20b and 35.3-20c can be verified using pari/gp as shown in figure 35.3-D.

35.3.3.5 An expression for a sum of products

Define Z_n as

$$Z_n \quad := \quad z_1 + z_1\,z_2 + z_1\,z_2\,z_3 + z_1\,z_2\,z_3\,z_4 + \ldots \;=\; \sum_{k=1}^{n} \prod_{i=1}^{k} z_i \tag{35.3-21}$$

Then Z_∞ has the continued fraction

$$Z_\infty \quad = \quad \cfrac{z_1}{1 - \cfrac{z_2}{1 + z_2 - \cfrac{z_3}{1 + z_3 - \cfrac{z_4}{1 + z_4 - \cfrac{z_5}{1 + z_5 - \cdots}}}}} \tag{35.3-22}$$

That is, $Z = K(a, b)$ where

$$a \quad = \quad [0,\ 1,\ z_2 + 1,\ z_3 + 1,\ z_4 + 1,\ z_5 + 1,\ z_6 + 1,\ \ldots] \tag{35.3-23a}$$

$$b \quad = \quad [1,\ z_1,\ -z_2,\quad -z_3,\quad -z_4,\quad -z_5,\quad -z_6,\quad \ldots] \tag{35.3-23b}$$

For the n-th convergent P_n/Q_n one has $Q_n = 1$ and $P_n = Z_n$.

To convert the power series of a hypergeometric function (see section 35.2 on page 685)

$$F\left(\left.\begin{matrix} a_1,\, a_2,\, \ldots,\, a_u \\ b_1,\, b_2,\, \ldots,\, b_v \end{matrix}\right| z\right) \tag{35.3-24}$$

into a continued fraction, set $z_1 = 1$, and for $k \geq 1$ set

$$z_{k+1} \quad = \quad \frac{z \prod_{j=1}^{u} (a_j + k)}{k \prod_{j=1}^{v} (b_j + k)} \tag{35.3-25}$$

An implementation is

```
1   hyper2cf(va, vb, n, z='z)=
2   \\ convert hypergeom(va,vb,z) to continued fraction
3   {
4       local(cfa, cfb, m);
5       n += 2;
6       cfa = vector(n);
7       cfb = vector(n);
8       cfa[1] = 0;    cfa[2] = 1;
9       cfb[1] = 1;    cfb[2] = 1;
10      for (k=3, n,
11          m = 1/(k-2);  \\ hidden lower parameter 1:  (n-2) == 1+(n-3)
12          m *= prod(j=1, #va, va[j]+(k-3));  \\ upper parameters
13          m /= prod(j=1, #vb, vb[j]+(k-3));  \\ lower parameters
14          m *= z;  \\ argument
15          cfa[k]=(m+1);
16          cfb[k]=-m;
17      );
18      return( [cfa, cfb] );
19  }
```

We convert $\log(1-z)/z = F\left(\begin{matrix}1,1\\2\end{matrix}\middle| z\right)$ to a continued fraction and check the result:

```
? N=7;
? va=[1,1];vb=[2];
? t=hyper2cf(va,vb,N);
? cfa=t[1]
  [0, 1, 1/2*z + 1, 2/3*z + 1, 3/4*z + 1, 4/5*z + 1, 5/6*z + 1, 6/7*z + 1, 7/8*z + 1]
? cfb=t[2]
  [1, 1, -1/2*z, -2/3*z, -3/4*z, -4/5*z, -5/6*z, -6/7*z, -7/8*z]
? t=cfab2pq(cfa,cfb)
  [1/8*z^7 + 1/7*z^6 + 1/6*z^5 + 1/5*z^4 + 1/4*z^3 + 1/3*z^2 + 1/2*z + 1, 1]
? s1=t[1]/t[2]+O(z^N)
  1 + 1/2*z + 1/3*z^2 + 1/4*z^3 + 1/5*z^4 + 1/6*z^5 + 1/7*z^6 + O(z^7)
? s2=hypergeom(va,vb,z,N)+O(z^N)
  1 + 1/2*z + 1/3*z^2 + 1/4*z^3 + 1/5*z^4 + 1/6*z^5 + 1/7*z^6 + O(z^7)
```

For further information on continued fractions see [208] and [133]. An in depth treatment is [188].

Chapter 36

Synthetic Iterations *

It is easy to construct arbitrary many iterations that converge super-linearly. Guided by some special constants that in base 2 can be obtained by recursive constructions we build iterations that allow the computation of the constant in a base independent manner. The iterations lead to functions that typically cannot be identified in terms of known (named) functions. Some of the functions can be expressed as infinite sums or products. For the constructions with repeated string substitutions see chapter 16.

36.1 A variation of the iteration for the inverse

We start with the product form for the most simple iteration, the one for the inverse:

$$\begin{aligned} I(y) &:= \frac{1}{1-y} && \text{(36.1-1a)}\\ &= 1 + y + y^2 + y^3 + y^4 + \dots && \text{(36.1-1b)}\\ &= (1+y)\,(1+y^2)\,(1+y^4)\,(1+y^8)\,\dots\,(1+y^{2^k})\,\dots && \text{(36.1-1c)}\\ &= (1+Y_0)\,(1+Y_1)\,(1+Y_2)\,(1+Y_3)\,\dots\,(1+Y_k)\,\dots && \text{(36.1-1d)}\\ &\quad \text{where} \quad Y_0 = y,\ Y_{k+1} = Y_k^2 \end{aligned}$$

We now modify the signs in the infinite product:

$$\begin{aligned} J(y) &:= (1-y)\,(1-y^2)\,(1-y^4)\,(1-y^8)\,\dots\,(1-y^{2^k})\,\dots && \text{(36.1-2a)}\\ &= 1 - y - y^2 + y^3 - y^4 + y^5 + y^6 - y^7 - y^8 \pm \dots && \text{(36.1-2b)}\\ &= (1-Y_0)\,(1-Y_1)\,(1-Y_2)\,(1-Y_3)\,\dots\,(1-Y_k)\,\dots && \text{(36.1-2c)}\\ &\quad \text{where} \quad Y_0 = y,\ Y_{k+1} = Y_k^2 \end{aligned}$$

The value of the n-th coefficient equals plus one if the parity of n is zero, else minus one (sequence A106400 of [240], the *Thue-Morse sequence*). The function J can be implemented as

```
fj(y,N=5)=
{
    local(r);
    r = 1;
    for (k=1, N,
        r -= r*y;
        y *= y;
    );
    return(r);
}
```

Replacing the minus by a plus gives the implementation for the function I.

A related constant is the *parity number* (or *Prouhet-Thue-Morse constant*):

$$\begin{aligned} P &= 0.412454033640107597783361368258455283089478374455769557 5\ldots \qquad (36.1\text{-}3)\\ [\text{base } 2] &= 0.0110,1001,1001,0110,1001,0110,0110,1001,1001,0110,0110,1001,\ldots\\ [\text{base } 16] &= 0.6996,9669,9669,6996,9669,6996,6996,9669,9669,6996,6996,9669,\ldots\\ [\text{CF}] &= [0,2,2,2,1,4,3,5,2,1,4,2,1,5,44,1,4,1,2,4,1,1,1,5,14,1,50,15,5,1,1,1,4,2,1,\ldots] \end{aligned}$$

The sequence of zeros and ones in the binary expansions is entry A010060 of [240]. The constant P can be computed defining

$$K(y) = \frac{[I(y)-J(y)]}{2} = y+y^2+y^4+y^7+y^8+y^{11}+y^{13}+y^{14}+y^{16}+\ldots \qquad (36.1\text{-}4)$$

Then

$$P = \frac{1}{2}K\left(\frac{1}{2}\right) = \frac{1}{2}\frac{\left[I(\frac{1}{2})-J(\frac{1}{2})\right]}{2} = \frac{1}{2}-\frac{1}{4}J\left(\frac{1}{2}\right) \qquad (36.1\text{-}5)$$

Thereby (see also [31], item 125),

$$2-4\,P = J\left(\frac{1}{2}\right) = \prod_{k=0}^{\infty}\left(1-\frac{1}{2^{2^k}}\right) = 0.350183865439569608866554526966178\ldots \qquad (36.1\text{-}6)$$

The sequence of bits of the parity number can also be obtained by starting with a single 0 and repeated application of the substitution rules $0 \to 01$ and $1 \to 10$.

The following relations are direct consequences of the definitions of the functions I and J:

$$I(y)\,I(-y) = I(y^2) \qquad (36.1\text{-}7a)$$

$$I(y) = \frac{J(y^2)}{J(y)} \qquad (36.1\text{-}7b)$$

$$I(-y) = \frac{1-y}{1+y}\,I(y) \qquad (36.1\text{-}7c)$$

$$J(-y) = \frac{1+y}{1-y}\,J(y) \qquad (36.1\text{-}7d)$$

We have

$$I(y) = 1+\sum_{k=0}^{\infty}\left[y^{2^k}\prod_{j=0}^{k-1}\left(1+y^{2^j}\right)\right] \qquad (36.1\text{-}8a)$$

$$J(y) = 1-\sum_{k=0}^{\infty}\left[y^{2^k}\prod_{j=0}^{k-1}\left(1-y^{2^j}\right)\right] \qquad (36.1\text{-}8b)$$

A functional equation for K is

$$K(y) = (1-y)\,K(y^2)+\frac{y}{1-y^2} \qquad (36.1\text{-}9)$$

It is solved by

$$K(y) = \sum_{k=0}^{\infty}\left[\frac{y^{2^k}}{1-y^{2^{k+1}}}\prod_{j=0}^{k-1}\left(1-y^{2^j}\right)\right] \qquad (36.1\text{-}10)$$

For the inverse of J we have

$$\frac{1}{J(y)} = 1 + y + 2y^2 + 2y^3 + 4y^4 + 4y^5 + 6y^6 + 6y^7 + 10y^8 + 10y^9 + 14y^{10} + \ldots \quad (36.1\text{-}11a)$$

$$= \left[(1-y)\,(1-y^2)\,(1-y^4)\,(1-y^8)\,\ldots\right]^{-1} = \prod_{k=0}^{\infty} I(y^{2^k}) \quad (36.1\text{-}11b)$$

$$= (1-y)\prod_{k=0}^{\infty}\frac{1+y^{2^k}}{1-y^{2^k}} \quad (36.1\text{-}11c)$$

$$= (1+y)\,(1+y^2)^2\,(1+y^4)^3\,(1+y^8)^4\,(1+y^{16})^5\,\ldots\,(1+y^{2^k})^{k+1}\,\ldots \quad (36.1\text{-}11d)$$

The sequence of coefficients of the even powers of x in relation 36.1-11a is

```
1, 2, 4, 6, 10, 14, 20, 26, 36, 46, 60, 74, 94, 114, 140, 166, 202, ...
```

This is entry A000123 of [240], the number of binary partitions of the even numbers. The sequence $\frac{1}{2}\,[2, 4, 6, 10, 14, \ldots]$ modulo two equals the period-doubling sequence, see section 36.5 on page 722.

Relation 36.1-11d tells us that the number of partitions of n into parts 2^k equals the number of partitions of n into at most one 1, two 2, three 4, ..., $k+1$ parts 2^k. The same relation can be used for a divisionless algorithm for the computation of $1/J$:

```
1   binpart(y,N=5)=
2   {
3       local(r);
4       r = 1;
5       for (k=1, N,
6           for (j=1, k,  r += r*y; );
7           y *= y;
8       );
9       return(r);
10  }
```

The generating function $1 + 2\,y + 4\,y^2 + 6\,y^3 + 10\,y^4 + 14\,y^5 + 20\,y^6 + \ldots$ equals

$$\frac{I(y)}{J(y)} = (1+y)^2\,(1+y^2)^3\,(1+y^4)^4\,(1+y^8)^5\,(1+y^{16})^6\,\ldots\,(1+y^{2^k})^{k+2}\,\ldots \quad (36.1\text{-}12)$$

It can be computed via (note the change in the inner loop)

```
1   binpart2(y,N=5)=
2   {
3       local(r);
4       r = 1;
5       for (k=1, N,
6           for (j=1, k+1,  r += r*y; );
7           y *= y;
8       );
9       return(r);
10  }
```

For $I(y)$ we have

$$I(y) = \sum_{k=0}^{\infty}\frac{y^{2^k-1}}{1-y^{2^{k+1}}} = \sum_{k=0}^{\infty}\frac{2^k\,y^{2^k-1}}{1+y^{2^k}} \quad (36.1\text{-}13a)$$

Integration gives

$$-\log(1-y) = \sum_{k=0}^{\infty}\frac{1}{2^{k+1}}\,\log\left(\frac{1+y^{2^k}}{1-y^{2^k}}\right) = \sum_{k=0}^{\infty}\log\left(1+y^{2^k}\right) \quad (36.1\text{-}14a)$$

For the derivative of J we have

$$J'(y) = -J(y)\sum_{k=0}^{\infty}\frac{2^k\,y^{2^k-1}}{1-y^{2^k}} \quad (36.1\text{-}15)$$

The following functional equations hold for $I(y)$:

$$0 = B - 2\,A\,B + A^2 \quad \text{where} \quad A = I(y), \quad B = I(y^2) \tag{36.1-16a}$$
$$0 = B - 2\,A\,B - A^2 + 2\,A^2\,B \quad \text{where} \quad A = I(-y), \quad B = I(-y^2) \tag{36.1-16b}$$
$$0 = B - 3\,A\,B + 3\,A^2\,B - A^3 \quad \text{where} \quad A = I(y), \quad B = I(y^3) \tag{36.1-16c}$$
$$0 = B - 5\,A\,B + 10\,A^2\,B - 10\,A^3\,B + 5\,A^4\,B - A^5 \quad \text{where} \quad A = I(y), \quad B = I(y^5) \tag{36.1-16d}$$
$$0 = B\left[(1-A)^k - (-A)^k\right] + (-A)^k \quad \text{where} \quad A = I(y), \quad B = I(y^k) \tag{36.1-16e}$$

From the functional equation for $K(y)$ (relation 36.1-9), the definition of $K(y)$, and relation 36.1-7b one can derive the following relation for $J(y)$:

$$0 = J_2^3 - 2\,J_4\,J_2\,J_1 + J_4\,J_1^2 \quad \text{where} \quad J_1 = J(y), \quad J_2 = J(y^2), \quad J_4 = J(y^4) \tag{36.1-17}$$

This relation is given in entry A106400 of [240], together with

$$0 = J_6\,J_1^3 - 3\,J_6\,J_2\,J_1^2 + 3\,J_6\,J_2^2\,J_1 - J_3\,J_2^3 \tag{36.1-18}$$

where $J_k = J(y^k)$. Relations between J_1, J_2, J_k, and J_{2k} can be derived from relation 36.1-16e by replacing $I(y)$ by $J(y^2)/J(y)$. For example, $k = 5$ gives

$$0 = J_{10}\,J_1^5 - 5\,J_{10}\,J_2\,J_1^4 + 10\,J_{10}\,J_2^2\,J_1^3 - 10\,J_{10}\,J_2^3\,J_1^2 + 5\,J_{10}\,J_2^4\,J_1 - J_5\,J_2^5 \tag{36.1-19}$$

The Komornik-Loreti constant

One has $K(1/\beta) = 1$ for

$$\frac{1}{\beta} = 0.55952455849672652513220976515743228583107647896866030 76\ldots \tag{36.1-20a}$$
$$[\text{base } 2] = 0.10001111001111010000000001100000000011010110001000100 10110\ldots$$
$$\beta = 1.78723165018296593301327489033700838533793140296181099 7\ldots \tag{36.1-20b}$$
$$[\text{base } 2] = 1.11001001100010000000001101101111111010010111010110100 00\ldots$$
$$[\text{CF}] = [1, 1, 3, 1, 2, 3, 188, 1, 12, 1, 1, 22, 33, 1, 10, 1, 1, 7, 1, 9, 1, 1, 20, 2, 15, 1, \ldots]$$

The constant β is the smallest real number in the interval $(1, 2)$ so that 1 has a unique expansion of the form $\sum_{n=1}^{\infty} \delta_n\,\beta^{-n}$ where $\delta_n \in \{0, 1\}$. It is called the *Komornik-Loreti constant* (see [8]). The fact that $\delta_n = 1$ exactly where the Thue-Morse sequence equals 1 was used for the computation of β: one solves $K(y) = 1$ for y. The transcendence of β is proved (using the fact that $J(y)$ is transcendental for algebraic y) in [9].

Third order variants

Variations of the third order iteration for the inverse

$$I(y) := \frac{1}{1-y} = 1 + y + y^2 + y^3 + y^4 + \ldots \tag{36.1-21a}$$
$$= (1 + y + y^2)\,(1 + y^3 + y^6)\,(1 + y^9 + y^{18})\,\ldots\,(1 + y^{3^k} + y^{2\cdot 3^k})\,\ldots \tag{36.1-21b}$$
$$= (1 + Y_0 + Y_0^2)\,(1 + Y_1 + Y_1^2)\,(1 + Y_2 + Y_2^2)\,\ldots\,(1 + Y_k + Y_k^2)\,\ldots \quad \text{where} \quad Y_0 = y,\ Y_{k+1} = Y_k^3 \tag{36.1-21c}$$

lead to series related to the base-3 analogue of the parity. The most simple example may be

$$T(y) = (1 + Y_0 - Y_0^2)\,(1 + Y_1 - Y_1^2)\,(1 + Y_2 - Y_2^2)\,\ldots\,(1 + Y_k - Y_k^2)\,\ldots \tag{36.1-22a}$$
$$= 1 + y - y^2 + y^3 + y^4 - y^5 - y^6 - y^7 + y^8 + y^9 + y^{10} - y^{11} + y^{12} \pm \ldots \tag{36.1-22b}$$

The sign of the n-th coefficient is the parity of the number of twos in the radix-3 expansion of n. We have

$$\frac{1}{2}\left(I(y)-T(y)\right) = y^2+y^5+y^6+y^7+y^{11}+y^{14}+y^{15}+y^{16}+y^{18}+y^{19}+y^{21}+\dots \quad (36.1\text{-}23a)$$

$$\frac{1}{4}\left(I\left(\frac{1}{2}\right)-T\left(\frac{1}{2}\right)\right) = 0.152644523625407582531924921475791679311504514871489254 8\dots \quad (36.1\text{-}23b)$$

$$[\text{base } 2] = 0.00100111000100111011011000100100111000100111011011000 11\dots$$

$$[\text{CF}] = [0,6,1,1,4,2,1,1,2,4,1,1,4,1,4,2,1,1,1,2,1,18,3,24,1,6,1,3,\dots]$$

The sequence of zeros and ones in the binary expansion is entry A064990 of [240], the *Mephisto Waltz sequence*.

36.2 An iteration related to the Thue constant

We construct a sequence of zeros and ones that can be generated by starting with a single 0 and repeated application of the substitution rules $0 \to 111$ and $1 \to 110$. The evolution starting with a single zero is:

```
T0 =  0
T1 =  111
T2 =  110110110
   == 3 times 11.0
T3 =  110110111110110111110110111
   == 3 times 110110.111
T4 =  110110111110110111110110110110111101101111101101101101101111101101111101101110
   == 3 times 110110111110110111.110110110
T --> 110110111110110111110110110110111110110111110110110110111110110111110110111101...
```

The crucial observation is that is T_n is three time repeated the string $T'_{n-1}.T_{n-2}$ where T'_k consists of the first and second third of T_k. The length of the n-th string is 3^n. Let $T(y)$ be the function whose power series corresponds to the string T_∞:

$$T(y) = 1+y+y^3+y^4+y^6+y^7+y^8+y^9+y^{10}+y^{12}+y^{13}+y^{15}+y^{16}+y^{17}+y^{18}+\dots \quad (36.2\text{-}1)$$

It can be computed by the iteration

$$L_0 = 0, \qquad A_0 = 1+y, \qquad B_0 = y^2, \qquad Y_0 = y \quad (36.2\text{-}2a)$$

$$R_n = A_n + y^2 L_n \quad (36.2\text{-}2b)$$

$$L_{n+1} = A_n + B_n \quad (36.2\text{-}2c)$$

$$Y_{n+1} = Y_n^3 \quad (36.2\text{-}2d)$$

$$A_{n+1} = R_n\left(1+Y_{n+1}\right) \quad \to T(y) \quad (36.2\text{-}2e)$$

$$B_{n+1} = R_n Y_{n+1}^2 \quad (36.2\text{-}2f)$$

The implementation is slightly tricky:

```
1  th(y, N=5)=
2  {
3      local(L, R, A, B, y2, y3, t);
4      /* correct up to order 3^(N+1)-1 */
5      L=0;
6      A=1+y;  B=y^2; /* R = A.B */
7      for(k=1, N,
8        /* (L, A.B) --> (A.B,  A.L.A.L . A.L) */
9        y2 = y^2;
10       R = A + y2*L; /* A.L */
11       L = A + B;   /* next L = A.B */
12       y3 = y * y2;
13       B = R * (y3*y3);
14       A = R * (1+y3);    /* next A = A.L.A.L  */
15       y = y3;
16     );
17     return( A + B )
18 }
```

The *Thue constant* (which should be *Roth's constant*, see entry A014578 of [240]) can be computed as

$$\begin{aligned}
\frac{1}{2}T(\frac{1}{2}) &= 0.859099796854703104903572502841974202614239955559439087 4\ldots && (36.2\text{-}3)\\
[\text{base } 2] &= 0.110, 110, 111, 110, 110, 111, 110, 110, 110, 110, 110, 111, 110, 110, 111, 110, 110, \ldots\\
[\text{base } 8] &= 0.667, 667, 666, 667, 667, 666, 667, 667, 667, 667, 667, 666, 667, 667, 666, 667, 667, \ldots\\
[\text{CF}] &= [0, 1, 6, 10, 3, 2, 513, 1, 1, 2, 1, 4, 2, 6576668769, 1, 1, 4,\\
&\quad 1, 2, 2, 256, 1, 1, 2, 1, 2, 3, 1, 3, 3, 241785163922925834941235 3,\\
&\quad 1, 2, 3, 1, 3, 2, 1, 2, 1, 1, 256, 2, 2, 1, 4, 2, 3288334384,\\
&\quad 1, 1, 4, 1, 2, 2, 146, 2, 3, 3, 2, 1, 2, 1, 12, X, \ldots]
\end{aligned}$$

The term X in the continued fraction has 74 decimal digits. By construction the bits at positions n not divisible by three are one and otherwise the complement of the bit at position $n/3$. As a functional equation (see also section 36.5 on page 722):

$$y\,T(y) + y^3\,T(y^3) = \frac{y}{1-y} \qquad (36.2\text{-}4)$$

From this relation we can obtain a series for $T(y)$:

$$T(y) = \sum_{n=0}^{\infty} (-1)^n \frac{y^{3^n-1}}{1-y^{3^n}} \qquad (36.2\text{-}5)$$

36.3 An iteration related to the Golay-Rudin-Shapiro sequence

Define $Q(y)$ by

$$\begin{aligned}
L_0 &= 1, \qquad R_0 = y, \qquad Y_0 = y && (36.3\text{-}1a)\\
L_{n+1} &= L_n + R_n \quad \to Q(y) && (36.3\text{-}1b)\\
Y_{n+1} &= Y_n^2 && (36.3\text{-}1c)\\
R_{n+1} &= Y_{n+1}\,(L_n - R_n) && (36.3\text{-}1d)
\end{aligned}$$

then

$$Q(y) = 1 + y + y^2 - y^3 + y^4 + y^5 - y^6 + y^7 + y^8 + y^9 + y^{10} - y^{11} - y^{12} - y^{13} + y^{14} - y^{15} + \ldots \qquad (36.3\text{-}2)$$

The sequence of coefficients is the *Golay-Rudin-Shapiro sequence* (or *GRS sequence*, entry A020985 of [240], see also page 45). The constant

$$\begin{aligned}
Q &= 0.92924386959737885325397664472205076441287553952432552 22\ldots && (36.3\text{-}3)\\
[\text{base } 2] &= 0.1110, 1101, 1110, 0010, 1110, 1101, 0001, 1101, 1110, 1101, 1110, 0010, \ldots\\
[\text{base } 16] &= 0.ede2, ed1d, ede2, 12e2, ede2, ed1d, 121d, ed1d, ede2, ed1d, ede2, 12e2, \ldots\\
[\text{CF}] &= [0, 1, 13, 7, 1, 1, 15, 4, 1, 3, 1, 2, 2, 1000, 12, 2, 1, 6, 1, 1, 1, 1, 1, 1, 8, 2, 1, 1, 2, 4, 1, 1, 3, \ldots]
\end{aligned}$$

can be computed as

$$Q = \frac{1 + \frac{1}{2}\,Q(\frac{1}{2})}{2} \qquad (36.3\text{-}4)$$

The implementation using pari/gp:

```
1   qq(y, N=8)=
2   {
3       local(L, R, Lp, Rp);
4       /* correct up to order 2**(N+1) */
5       L=1;  R=y;
6       for(k=0,N, Lp=L+R; y*=y; Rp=y*(L-R); L=Lp; R=Rp);
7       return( L + R )
8   }
```

The following functional relations hold for Q:

$$Q(y^2) = \frac{Q(y)+Q(-y)}{2} \tag{36.3-5a}$$

$$Q(y) = Q(y^2) + yQ(-y^2) \tag{36.3-5b}$$

$$Q(-y) = Q(y^2) - yQ(-y^2) \tag{36.3-5c}$$

Combining the latter two gives

$$Q(y) = (1+y)\,Q(y^4) + (y^2-y^3)\,Q(-y^4) \tag{36.3-6}$$

Counting zeros and ones in the binary expansion of Q

The number of ones and zeros in the first 4^k bits of the constant Q can be computed via a string-substitution engine (see chapter 16 on page 349). The hexadecimal expansion can be obtained as:

```
Number of symbols = 4
Start: e
Rules:
  e --> ed
  d --> e2
  2 --> 1d
  1 --> 12
-------------
0:   (#=1)
  e
1:   (#=2)
  ed
2:   (#=4)
  ede2
3:   (#=8)
  ede2ed1d
4:   (#=16)
  ede2ed1dede212e2
5:   (#=32)
  ede2ed1dede212e2ede2ed1d121ded1d
6:   (#=64)
  ede2ed1dede212e2ede2ed1d121ded1dede2ed1dede212e2121d12e2ede212e2
```

A few lines of pari/gp code count the occurrences of the symbols (and thereby of zeros and ones):

```
1   /*   e --> ed  ;  d --> e2  ;  2 --> 1d  ;  1 --> 12    */
2   /*   e  d      ;  e     2   ;     d     1;        2  1; */
3   mg= [1, 1, 0, 0;  1, 0, 1, 0;  0, 1, 0, 1;  0, 0, 1, 1];
4   mg=mattranspose(mg)
5   { for (k=0, 40,
6       print1( k, ":  " );
7       mm=mg^k;
8       mv = mm*[1,0,0,0]~;
9       t =  sum(i=1,4, mv[i]);
10      /* e and d have three ones and one zero */
11      /* 1 and 2 have one one and three zeros */
12      n0 = 3*(mv[3]+mv[4]) + (mv[1]+mv[1]); /* # of zeros */
13      n1 = 3*(mv[1]+mv[2]) + (mv[3]+mv[4]); /* # of ones  */
14      print( t, "   ",  mv~,
15       "   #0=",  n0, " #1=", n1, "  diff=", n1-n0, "  #1/#0=", 1.0*n1/n0 );
16    ) }
```

We obtain the data shown in figure 36.3-A. The sequence of the numbers of ones is entry A005418 of [240]. It is identical to the sequence of numbers of equivalence classes obtained by identifying bit-strings that are mutual reverses or complements, see section 3.8.1.5 on page 138.

```
                #e   #d   #2   #1
 0:      1   [  1,   0,   0,   0]   #0=   2   #1=   3  diff= 1  #1/#0=1.5000000
 1:      2   [  1,   1,   0,   0]   #0=   2   #1=   6  diff= 4  #1/#0=3.0000000
 2:      4   [  2,   1,   1,   0]   #0=   7   #1=  10  diff= 3  #1/#0=1.4285714
 3:      8   [  3,   3,   1,   1]   #0=  12   #1=  20  diff= 8  #1/#0=1.6666666
 4:     16   [  6,   4,   4,   2]   #0=  30   #1=  36  diff= 6  #1/#0=1.2000000
 5:     32   [ 10,  10,   6,   6]   #0=  56   #1=  72  diff=16  #1/#0=1.2857142
 6:     64   [ 20,  16,  16,  12]   #0= 124   #1= 136  diff=12  #1/#0=1.0967741
 7:    128   [ 36,  36,  28,  28]   #0= 240   #1= 272  diff=32  #1/#0=1.1333333
 8:    256   [ 72,  64,  64,  56]   #0= 504   #1= 528  diff=24  #1/#0=1.0476190
 9:    512   [136, 136, 120, 120]   #0= 992   #1=1056  diff=64  #1/#0=1.0645161
10:   1024   [272, 256, 256, 240]   #0=2032   #1=2080  diff=48  #1/#0=1.0236220
  [--snip--]
40: 1099511627776   [274878431232, 274877906944, 274877906944, 274877382656]  \
        #0=2199022731264 #1=2199024304128  diff=1572864  #1/#0=1.00000071525
```

Figure 36.3-A: Number of symbols, zeros and ones with the n-th step of the string substitution engine for the GRS sequence. For long strings the ratio of the number of zeros and ones approaches one.

36.4 Iterations related to the ruler function

The *ruler function* $r(n)$ can be defined to be the highest exponent e so that 2^e divides n. Here we consider the function that equals $r(n)+1$ for $n \neq 0$ and zero for $n = 0$. The partial sequences up to indices $2^n - 1$ are

```
R0 = 0
R1 = 0 1
R2 = 0 1 2 1
R3 = 0 1 2 1 3 1 2 1
R4 = 0 1 2 1 3 1 2 1 4 1 2 1 3 1 2 1
R5 = 0 1 2 1 3 1 2 1 4 1 2 1 3 1 2 1 5 1 2 1 3 1 2 1 4 1 2 1 3 1 2 1
R6 = 0 1 2 1 3 1 2 1 4 1 2 1 3 1 2 1 5 1 2 1 3 1 2 1 4 1 2 1 3 1 2 1 6 1 ...
```

The limiting sequence is entry A001511 of [240]. Observe that $R_n = R_{n-1}.(R_{n-1} + [n, 0, 0, \ldots, 0])$. Define the function $R(y)$ as the limit of the iteration

$$R_1 = y, \qquad Y_1 = y \tag{36.4-1a}$$

$$Y_{n+1} = Y_n^2 \tag{36.4-1b}$$

$$R_{n+1} = R_n + Y_{n+1}\,[R_n + (1+n)] \quad \to R(y) \tag{36.4-1c}$$

Implementation in pari/gp:

```
1  r2(y, N=11)=
2  { /* correct to order = 2^N-1 */
3      local(A);
4      A=y;
5      for(k=2, N,  y *= y;  A += y*(A + k); );
6      return( A );
7  }
```

If one replaces in the statement `A += y*(A + k);` by `A += y*(A + 1);` then the iteration computes $\frac{y}{1-y}$. For the function R we have

$$R\left(\frac{1}{q^2}\right) = \frac{1}{2q}\left[(q-1)\,R\left(\frac{1}{q}\right) + (q+1)\,R\left(-\frac{1}{q}\right)\right] \tag{36.4-2a}$$

$$R\left(\frac{1}{q}\right) = R\left(\frac{1}{q^2}\right) + \frac{1}{q-1} \tag{36.4-2b}$$

$$R(y) = R(y^2) + \frac{y}{1-y} = R(y^4) + \frac{y}{1-y} + \frac{y^2}{1-y^2} = \ldots \tag{36.4-2c}$$

and so

$$R(y) = \sum_{n=0}^{\infty} \frac{y^{2^n}}{1-y^{2^n}} \tag{36.4-3}$$

One further has

$$R(y) = \sum_{n=0}^{\infty} (1+n)\,\frac{y^{2^n}}{1-y^{2^{n+1}}} \tag{36.4-4}$$

Define the *ruler constant* as $R := R(1/2)/2$, then

$$\begin{aligned} R &= 0.701968413941089160288103037068604677268819380760945033 7\ldots \qquad (36.4\text{-}5)\\ [\text{base } 2] &= 0.101100111011010000110011101101001011001110110100001100 11\ldots \\ [\text{CF}] &= [0,1,2,2,1,4,2,1,1,1,2,2,1,3,3,4,5,6,1,5,1,1,9,49,1,8,1,1,5,1,\\ &\quad 6,5,1,3,3,1,2,4,3,1,2,4,2,1,1,3,1,9,1,11,18,2,4,5,1,3,2,25,9,\\ &\quad 2,3,1,2,3,1,9,1,2,8,1,3,4,1,1,1,1,31,1,1,6,1,13,1,1,14,1,6,1,\ldots] \end{aligned}$$

We will now compute the function

$$P(y) \;=\; \sum_{n=0}^{\infty} \frac{y^{2^n}}{1+y^{2^n}} \;=\; \sum_{n=0}^{\infty} (1-n)\,\frac{y^{2^n}}{1-y^{2^{n+1}}} \qquad (36.4\text{-}6)$$

The partial sequences of coefficients of the Taylor expansion are

```
P0 = 0
P1 = 0 1
P2 = 0 1 0 1
P3 = 0 1 0 1 -1 1 0 1
P4 = 0 1 0 1 -1 1 0 1 -2 1 0 1 -1 1 0 1
P5 = 0 1 0 1 -1 1 0 1 -2 1 0 1 -1 1 0 1 -3 1 0 1 -1 1 0 1 -2 1 0 1 -1 1 0 1
P6 = 0 1 0 1 -1 1 0 1 -2 1 0 1 -1 1 0 1 -3 1 0 1 -1 1 0 1 -2 1 0 1 -1 1 0 1 -4 1...
```

Observe that $P_n = P_{n-1}.(P_{n-1} - [n-2, 0, 0, \ldots, 0])$. Compute $P(y)$ by

$$\begin{aligned} P_1 &= y, \qquad Y_1 = y && (36.4\text{-}7a)\\ Y_{n+1} &= Y_n^2 && (36.4\text{-}7b)\\ P_{n+1} &= P_n + Y_{n+1}\,[P_n + (1-n)] \quad \to P(y) && (36.4\text{-}7c) \end{aligned}$$

Implementation in pari/gp:

```
1  p2(y, N=11)=
2  { /* correct to order = 2^N-1 */
3      local(A);
4      A=y;
5      for(k=2, N,  y *= y;  A += y*(A - (k-2)); );
6      return( A );
7  }
```

One finds for P:

$$\begin{aligned} P\left(\frac{1}{q^2}\right) &= \frac{1}{2q}\left[(q+1)\,P\left(\frac{1}{q}\right) + (q-1)\,P\left(-\frac{1}{q}\right)\right] && (36.4\text{-}8a)\\ P\left(\frac{1}{q}\right) &= P\left(\frac{1}{q^2}\right) + \frac{1}{q+1} && (36.4\text{-}8b) \end{aligned}$$

Relations that involve both P and R are

$$\begin{aligned} P\left(\frac{1}{q^2}\right) &= \frac{1}{2}\left[\frac{q+3}{q+1}\,P\left(\frac{1}{q}\right) - \frac{q-1}{q+1}\,R\left(\frac{1}{q}\right)\right] && (36.4\text{-}8c)\\ R\left(\frac{1}{q}\right) &= \frac{1}{2}\left[(q+1)\,P\left(\frac{1}{q^2}\right) + (q+3)\,R\left(\frac{1}{q^2}\right)\right] && (36.4\text{-}8d)\\ P\left(\frac{1}{q}\right) &= \frac{1}{2}\left[(q+1)\,P\left(\frac{1}{q^2}\right) + (q-1)\,R\left(\frac{1}{q^2}\right)\right] && (36.4\text{-}8e) \end{aligned}$$

$$\begin{aligned} \frac{R(y)+P(y)}{2} &= \frac{y}{1-y} \;=\; \sum_{n=1}^{\infty} y^n && (36.4\text{-}8f)\\ \frac{R(y)-P(y)}{2} &= R(y^2) && (36.4\text{-}8g)\\ R^2(y) - P^2(y) &= \frac{4\,y}{1-y}\,R(y^2) && (36.4\text{-}8h) \end{aligned}$$

36.5 An iteration related to the period-doubling sequence

Define

$$T(y) = \sum_{n=0}^{\infty} \frac{y^{2^n}}{1+(-1)^n\, y^{2^n}} = \sum_{n=0}^{\infty} (-1)^n \frac{y^{2^n}}{1-y^{2^n}} \tag{36.5-1a}$$

$$= y + y^3 + y^4 + y^5 + y^7 + y^9 + y^{11} + y^{12} + y^{13} + y^{15} + y^{16} + y^{17} + \ldots \tag{36.5-1b}$$

The function can be computed by the iteration

$$A_1 = 0, \qquad L_1 = y, \qquad R_1 = y, \qquad Y_1 = y \tag{36.5-2a}$$

$$A_{n+1} = L_n + Y_n\, R_n \quad \to T(y) \tag{36.5-2b}$$

$$L_{n+1} = L_n + Y_n\, A_n \tag{36.5-2c}$$

$$R_{n+1} = R_n + Y_n\, A_n \tag{36.5-2d}$$

$$Y_{n+1} = Y_n^2 \tag{36.5-2e}$$

Implementation in pari/gp:

```
1   t2(y, N=11)=
2   { /* correct to order = 2^N-1 */
3       local(A, L, R, t);
4       A=0;  L=y;  R=y;
5       for(k=2, N,
6           t = y*A;
7           A = L + y*R;
8           L += t;  R += t;
9           y *= y;
10      );
11      return( A );
12  }
```

The computed Taylor series can be obtained symbolically by starting with a single zero and applying the substitution rules $0 \to 11$ and $1 \to 10$. The evolution is:

```
T0 =  0
T1 =  1 1
T2 =  1 0 1 0
T3 =  1 0 1 1 1 0 1 1
T4 =  1 0 1 1 1 0 1 0 1 0 1 1 1 0 1 0
T5 =  1 0 1 1 1 0 1 0 1 0 1 1 1 0 1 1 1 0 1 1 1 0 1 0 1 0 1 1 1 0 1 1
T6 =  1 0 1 1 1 0 1 0 1 0 1 1 1 0 1 1 1 0 1 1 1 0 1 0 1 0 1 1 1 0 1 0 1 0 1 1 ...
```

Observe that $T_n = L(T_{n-1}).T_{n-2}.R(T_{n-1}).T_{n-2}$ where L and R denote the left and right half of their arguments. The limiting sequence is the so-called *period-doubling sequence*. It is entry A035263 of [240] where it is called the *first Feigenbaum symbolic sequence*. Define the *period-doubling constant* as $T := T(1/2)$, then

$$T = 0.729427023494948405709066206894052617060026944465854741 7\ldots \tag{36.5-3}$$

$$[\text{base } 2] = 0.10111010101110111011101010111010101110101011101110111010\ldots$$

$$[\text{CF}] = [0, 1, 2, 1, 2, 3, 2, 8, 1, 1, 1, 2, 1, 8, 6, 1, 2, 1, 2, 8, 1, 2, 2, 1, 1, 24, 2, 2, 2, 1,$$
$$8, 2, 1, 2, 1, 8, 6, 1, 2, 1, 2, 2, 1, 1, 1, 13, 1, 1, 8, 2, 13, 1, 1, 1, 2, 2, 1, 2, 1,$$
$$6, 8, 1, 2, 1, 2, 8, 1, 2, 2, 1, 1, 24, 2, 2, 2, 1, 8, 2, 1, 2, 1, 11, 1, 9, 2, 1, 116, \ldots]$$

The transcendence of this constant is proved in [159]. A functional equation for $T(y)$ is

$$T(y) + T(y^2) = \frac{y}{1-y} \tag{36.5-4}$$

The Taylor series of $T(y^2)$ has coefficients one where the series of $T(y)$ has coefficients zero.

```
             pile_0  pile_+  pile_-  moved    summary  direction
                                      disk              of move
        0:    1111    ....    ....    ...1    0 0 0 0
        1:    111.    ....    ...1    ...1    0 0 0 -    1
        2:    11..    ..1.    ...1    ..1.    0 0 + -    0
        3:    11..    ..11    ....    ...1    0 0 + +    1
        4:    1...    ..11    .1..    .1..    0 - + +    1
        5:    1..1    ..1.    .1..    ...1    0 - + 0    1
        6:    1..1    ....    .11.    ..1.    0 - - 0    0
        7:    1...    ....    .111    ...1    0 - - -    1
        8:    ....    1...    .111    1...    + - - -    0
        9:    ....    1..1    .11.    ...1    + - - +    1
       10:    ..1.    1..1    .1..    ..1.    + - 0 +    0
       11:    ..11    1...    .1..    ...1    + - 0 0    1
       12:    ..11    11..    ....    .1..    + + 0 0    1
       13:    ..1.    11..    ...1    ...1    + + 0 -    1
       14:    ....    111.    ...1    ..1.    + + + -    0
       15:    ....    1111    ....    ...1    + + + +    1
```

Figure 36.5-A: Solution of the towers of Hanoi puzzle for 4 disks. The rightmost column corresponds to the direction of the move made, it is the period-doubling sequence.

36.5.1 Connection to the towers of Hanoi puzzle

The *towers of Hanoi* puzzle consists of three piles and n disks of different size. The initial configuration is that all disks are on the leftmost pile ordered by size (smallest on top). The task is to move all disks to the rightmost pile by moving only one disk at a time and never putting a bigger disk on top of a smaller one.

The puzzle with n disks can be solved in $2^n - 1$ steps. Figure 36.5-A shows the solution for $n = 4$ [FXT: bits/hanoi-demo.cc]. Here the piles are represented as binary words. Note that with each move the lowest bit in one of the three words is moved to another word where it is again the lowest bit.

A simple solution can be obtained by observing that the disk moved with step $k = 1, \ldots, 2^n - 1$ corresponds to the lowest set bit in the binary representation of k and the index of the untouched pile changes by $+1$ mod 3 for n even and -1 mod 3 for n odd. The essential part of the implementation is

```
1   void
2   hanoi(ulong n)
3   {
4       ulong f[3];
5       f[0] = first_comb(n);  f[1] = 0;  f[2] = 0;  // Initial configuration
6
7       const int dr = (n&1 ? -1 : +1);  // == +1 (if n even),  else  == -1
8
9       // PRINT configuration
10      int u;  // index of tower untouched in current move
11      if ( dr<0 ) u=2;  else u=1;
12
13      ulong n2 = 1UL<<n;
14      for (ulong k=1; k<n2; ++k)
15      {
16          ulong s = lowest_one(k);
17
18          ulong j = 3;  while ( j-- )  f[j] ^= s;  // change all piles
19          f[u] ^= s;  // undo change for untouched pile
20
21          u += dr;
22          if ( u<0 ) u=2;  else if ( u>2 ) u=0;  // modulo 3
23
24          // PRINT configuration
25      }
26  }
```

Now with each step the transferred disk is moved by $+1$ or -1 position (modulo 3). The rightmost column in figure 36.5-A consists of zeros and ones corresponding to the direction of the move. It is the period-doubling sequence. A recursive algorithm for the towers of Hanoi puzzle can be given as [FXT: comb/hanoi-rec-demo.cc]:

```
1   ulong f[3];  // the three piles
2   void hanoi(int k, ulong A, ulong B, ulong C)
3   // Move k disks from pile A to pile C
```

```
4   {
5       if ( k==0 ) return;
6
7       // 1. move k-1 disks from pile A to pile B:
8       hanoi(k-1, A, C, B);
9
10      // 2. move disk k from pile A to pile C:
11      ulong b = 1UL << (k-1);
12      f[A] ^= b;
13      f[C] ^= b;
14
15      print_hanoi(b); // visit state
16
17      // 3. move k-1 disks from pile B to pile C:
18      hanoi(k-1, B, A, C);
19  }
```

The piles are represented by the binary words `f[A]`, `f[B]`, and `f[C]`, the variable `k` is the number of the disk moved. The routine is called as follows

```
1       ulong n = 5;
2       // Initial configuration:
3       f[0] = first_comb(n);  // n ones as lowest bits
4       f[1] = 0;  f[2] = 0;   // empty
5
6       // visit initial state
7
8       hanoi(n-1, 0, 1, 2);  // solve
```

36.5.2 Generalizations of the period-doubling sequence

```
? default(realprecision,85);
? tm(y,m,N=8)=sum(n=0,N,(-1)^n*(y^(m^n)/(1-y^(m^(n)))))
? for(m=2,9,print(m,": ",tm(1.0/10,m,4));print("   ",tm(1.0/10^m,m,4)))
    2: 1.111.1.1.111.111.111.1.1.111.111.111.1.1.111.111.111.1.1.111.111.111.1.1.111.111.111
       .1...1.1.1...1...1...1.1.1...1.1.1...1.1.1...1...1...1.1.1...1.1.1...1.1.1...1...1...
    3: 11.11.11111.11.11111.11.11.11.11.11111.11.11111.11.11.11.11.11111.11.11111.11.11111.1
       ..1..1.....1..1.....1..1..1..1..1.....1..1.....1..1..1..1..1.....1..1.....1..1.....1.
    4: 111.111.111.1111111.111.111.1111111.111.111.1111111.111.111.111.111.111.111.1111111.1
       ...1...1...1.......1...1...1.......1...1...1.......1...1...1...1...1...1...1.......1.
    5: 1111.1111.1111.1111.111111111.1111.1111.1111.111111111.1111.1111.1111.111111111.1111.
       ....1....1....1....1.........1....1....1....1.........1....1....1....1.........1....1
    6: 11111.11111.11111.11111.11111.1111111111.11111.11111.11111.11111.1111111111.11111.1
       .....1.....1.....1.....1.....1...........1.....1.....1.....1.....1...........1.....1.
    7: 111111.111111.111111.111111.111111.111111.1111111111111.111111.111111.111111.111111.1
       ......1......1......1......1......1......1.............1......1......1......1......1.
    8: 1111111.1111111.1111111.1111111.1111111.1111111.1111111.111111111111111.1111111.11111
       .......1.......1.......1.......1.......1.......1.......1...............1.......1.....
    9: 11111111.11111111.11111111.11111111.11111111.11111111.11111111.11111111.111111111111
       ........1........1........1........1........1........1........1........1............
```

Figure 36.5-B: Visual verification of the higher order analogues of the period-doubling sequence. In the output lines the leading '0.' was removed and all zeros were replaced by dots.

The functional equation for the period-doubling sequence, relation 36.5-4 can be generalized in several ways. For example, one can look for a function for which $F_3(y) + F_3(y^3) = \frac{y}{1-y}$. It is given by

$$F(y) = \sum_{k=0}^{\infty} (-1)^n \frac{y^{3^n}}{1-y^{3^n}} \tag{36.5-5a}$$

$$= z + z^2 + z^4 + z^5 + z^7 + z^8 + z^9 + z^{10} + z^{11} + z^{13} + z^{14} + z^{16} + z^{17} + \ldots \tag{36.5-5b}$$

We can compute a constant:

$$F_3(1/2) = 0.85909979685470310490357250284197420261423995555943908 74\ldots \tag{36.5-6}$$

But this is just the Thue constant, see section 36.2 on page 717. Large terms occur in the continued fraction expansion of this constant. Even larger terms occur in the continued fractions of $F_m(1/2)$ for

$m > 3$ (replace 3 by m in relation 36.5-5a). For example, the 45-th term of $F_5(1/2)$ has 565 digits. In contrast, the greatest of the first 1630 terms of the continued fraction of $F_2(1/2) = T(1/2)$ equals 288. Some sequences corresponding to the higher order analogues of the period-doubling sequence are shown in figure 36.5-B.

A different way to generalize is to search functions for which, for example, the following functional equation holds:

$$F(y) + F(y^2) + F(y^3) \quad = \quad \frac{y}{1-y} \tag{36.5-7}$$

The equation can be solved by writing $F(y) = y/(1-y) - F(y^2) - F(y^3)$ and using recursions that terminate when a prescribed order is reached.

```
1   F(z, R)=
2   {/* solve F(y) + F(y^2) + F(y^3) = y/(1-y) */
3       local(s, y);
4       y = z + R;
5       s = y/(1-y);
6       if ( y^2!=R,  s -= F(z^2, R) );
7       if ( y^3!=R,  s -= F(z^3, R) );
8       return(s);
9   }
```

We obtain

```
? N=55;
? default(seriesprecision,N);
? s=F(y,O(y^N))
  y + y^4 + y^5 + y^6 + y^7 + y^9 + y^11 - y^12 + y^13 + y^16 + y^17 -
  y^18 + y^19 + y^20 + y^23 + 2*y^24 + y^25 + y^28 + y^29 + y^30 + y^31 +
  y^35 + 3*y^36 + y^37 + y^41 + y^42 + y^43 + y^44 + y^45 + y^47 -
  2*y^48 + y^49 + y^52 + y^53 + 2*y^54 + O(y^55)
```

To verify that the function F actually satisfies the given functional equation we show the sequences of coefficients of the power series of $F(y)$, $F(y^2)$, $F(y^3)$ and their sum:

```
  F(y)=   [0, 1, 0, 0, 1, 1, 1, 1, 0, 1, 0, 1,-1, 1, 0, 0, 1, 1,-1, 1, 1, 0, 0, 1, 2, 1, ...]
  F(y^2)= [0, 0, 1, 0, 0, 0, 0, 0, 1, 0, 1, 0, 1, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0,-1, 0, ...]
  F(y^3)= [0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0, ...]
  sum=    [0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, ...]
```

$F(y) = y + y^4 + y^5 + y^7 + y^9 + y^{11} + y^{13} + y^{16} + y^{17} + y^{19} + y^{20} + y^{23} + y^{25} + y^{28} + y^{29} + \ldots$

```
F(y)=   .1..11.1.1.1.1..11.11..1.1..11.1...111...1.111.1.1..11.1...1.1.111.11..1.1
F(y^2)= ..1.....1.1...1...1...1...1.....1.1...1.1.....1...1.....1.1...1.......1.1.
F(y^3)= ...1........1..1.....1.....1.....1.....1........1..1.....1..1........1....
F(y^6)= ......1.................1.....1...........1...........1...........1.......
sum=    .1111111111111111111111111111111111111111111111111111111111111111111111111
```

Figure 36.5-C: Visual demonstration that $F(y)$ satisfies the functional equation $F(y)+F(y^2)+F(y^3)+F(y^6) = y/(1-y)$. Dots are used for zeros.

The Taylor series of $F(y)$ where $F(y) + F(y^2) + F(y^3) + F(y^6) = y/(1-y)$ contains only ones and zeros. One has

$$F(y) \quad = \quad \sum_{k=1}^{\infty} R(k)\, \frac{x^k}{1-x^k} \quad \text{where} \quad R(k) \;=\; \begin{cases} (-1)^{e_2+e_3} & \text{if} \quad k = 2^{e_2}\, 3^{e_3} \\ 0 & \text{else} \end{cases} \tag{36.5-8}$$

This is a Lambert series, it can be converted into a Taylor series by relation 35.1-20 on page 680. It turns out that with $k = u\, 2^{e_2}\, 3^{e_3}$ (where neither 2 nor 3 divides u), we have

$$r(k) \quad = \quad \begin{cases} 0 & \text{if either of } e_2 \text{ or } e_3 \text{ is odd} \\ 1 & \text{else} \quad \text{(i.e. both } e_2 \text{ and } e_3 \text{ are even)} \end{cases} \tag{36.5-9}$$

36.6 An iteration from substitution rules with sign

Substitution rules: 0 → L1, 1 → 10, and L → L0.

```
D0 =  1
D1 =  10
D2 =  10L1
D3 =  10L1L010
D4 =  10L1L010L01L10L1
D5 =  10L1L010L01L10L1L01L10L010L1L010
D --> 10L1L010L01L10L1L01L10L010L1L010L01L10L010L1L01L10L1L010L01L...
```

Identify L with -1 and observe that for $n > 1$ one has $D_n = D_{n-1}.(-D_{n-2}).D_{n-2}$. The length of the strings is 2^n. Define $D(y)$ by the iteration

$$\begin{aligned} L_0 &= 1, \qquad R_0 = 1(+0\,y), \qquad Y_0 = y && \text{(36.6-1a)}\\ L_{n+1} &= R_n && \text{(36.6-1b)}\\ R_{n+1} &= R_n + Y_n^2\,(-L_n + Y_n\,L_n) \quad \to D(y) && \text{(36.6-1c)}\\ Y_{n+1} &= Y_n^2 && \text{(36.6-1d)} \end{aligned}$$

Implementation in pari/gp:

```
1   dd(y, N=7)=
2   {
3       local(R, L, y2, t);
4       /* correct up to order 2^(N) */
5       L = 1;
6       R = 1 + 0*y;
7       for(k=1, N,
8         /* (L, R) --> (R,  R.(-L).L) */
9         y2 = y^2;
10        t = R;
11        R = R + y2*(-L + y*L);  /* R.(-L).L */
12        L = t;
13        y = y2;
14      );
15      return( R )
16  }
```

We have

$$D(y) = 1 - y^2 + y^3 - y^4 + y^6 - y^8 + y^{10} - y^{11} + y^{12} - y^{14} + y^{15} - y^{16} + y^{18} \pm \ldots \quad \text{(36.6-2)}$$

and (see relation 36.1-4 on page 714)

$$\frac{y}{1-y}\,D(y) = K(y) = y + y^2 + y^4 + y^7 + y^8 + y^{11} + \ldots \quad \text{(36.6-3)}$$

The coefficients are one where the Thue-Morse sequence equals minus one. Thereby the parity number can be computed as $P = \frac{1}{2}\,D\left(\frac{1}{2}\right)$. A functional equation for D is

$$D\,(y) = y\,\frac{1-y}{1+y}\,D\left(y^2\right) + \frac{1}{1+y} \quad \text{(36.6-4)}$$

36.7 Iterations related to the sum of digits

The sequence of the sum of binary digits of the natural numbers starting with zero (entry A035263 of [240]) can be constructed as follows:

```
S0 = 0
S1 = 0 1
S2 = 0 1 1 2
S3 = 0 1 1 2 1 2 2 3
S4 = 0 1 1 2 1 2 2 3 1 2 2 3 2 3 3 4
S5 = 0 1 1 2 1 2 2 3 1 2 2 3 2 3 3 4 1 2 2 3 2 3 3 4 2 3 3 4 3 4 4 5
```

Observe that $S_n = S_{n-1}.(S_{n-1} + I_{n-1})$ where I_n is a sequence of n ones and addition is element-wise.

Define $S(y)$ by

$$I_1 = 1, \qquad A_1 = y, \qquad Y_1 = y \tag{36.7-1a}$$

$$I_{n+1} = I_n\,(1+Y_n) = \sum_{k=0}^{2^n-1} y^k \tag{36.7-1b}$$

$$Y_{n+1} = Y_n^2 \tag{36.7-1c}$$

$$A_{n+1} = A_n + Y_{n+1}\,(I_{n+1}+A_n) \quad \to S(y) \tag{36.7-1d}$$

Implementation in pari/gp:

```
1   s2(y, N=7)=
2   {
3       local(in, A);
4       /* correct to order = 2^N-1 */
5       in = 1; /* 1+y+y^2+y^3+...+y^(2^k-1) */
6       A = y;
7       for(k=2, N,
8           in *= (1+y);
9           y *= y;
10          A += y*(in + A);
11      );
12      return( A );
13  }
```

The Taylor series is

$$S(y) = 0 + y + y^2 + 2y^3 + y^4 + 2y^5 + 2y^6 + 3y^7 + y^8 + 2y^9 + 2y^{10} + 3y^{11} + 2y^{12} + 3y^{13} + \ldots \tag{36.7-2}$$

Define the *sum-of-digits constant* as $S := S(1/2)/2$, then

$$S = 0.596063172117821679423793925862790645462361238478109932 6\ldots \tag{36.7-3}$$

$$[\text{base } 2] = 0.10011000100101111001100010010110100110001001011110011000 1\ldots$$

$$[\text{CF}] = [0, 1, 1, 2, 9, 1, 3, 5, 1, 2, 1, 1, 1, 1, 8, 2, 1, 1, 2, 1, 12, 19, 24, 1, 18, 12, 1, \ldots]$$

We have (see [247])

$$S(y) = \frac{1}{1-y}\sum_{k=0}^{\infty}\frac{y^{2^k}}{1+y^{2^k}} \tag{36.7-4}$$

and also

$$S(y) = \sum_{k=0}^{\infty}\left[\frac{y^{2^k}}{1-y^{2^k}}\prod_{j=0}^{k-1}\left[1+y^{2^j}\right]\right] \tag{36.7-5}$$

The last relation follows from the functional relation for S,

$$S(y) = (1+y)\,S(y^2) + \frac{y}{1-y^2} \tag{36.7-6}$$

It is of the form $F(y) = A(y)\,F(y^2) + B(y)$ where $A(y) = 1+y$ and $B(y) = y/(1-y^2)$ and has the solution

$$F(y) = \sum_{k=0}^{\infty}\left[B(y^{2^k})\prod_{j=0}^{k-1}\left[A(y^{2^j})\right]\right] \tag{36.7-7}$$

This can be seen by applying the functional equation several times:

$$F(y) = A(y)\,F(y^2) + B(y) \tag{36.7-8a}$$

$$= A(y)\,\left[A(y^2)\,F(y^4) + B(y^2)\right] + B(y) \tag{36.7-8b}$$

$$= A(y)\,\left[A(y^2)\,\left[A(y^4)\,F(y^8) + B(y^4)\right] + B(y^2)\right] + B(y) = \ldots \tag{36.7-8c}$$

$$= B(y) + A(y)\,B(y^2) + A(y)\,A(y^2)\,B(y^4) + \ldots \tag{36.7-8d}$$

Weighted sum of digits

Define $W(y)$ by

$$I_1 = \frac{1}{2}, \qquad A_1 = 1, \qquad Y_1 = y \tag{36.7-9a}$$

$$l_{n+1} = I_n \frac{(1+Y_n)}{2} = \frac{1}{2^n} \sum_{k=0}^{2^n-1} y^k \tag{36.7-9b}$$

$$Y_{n+1} = Y_n^2 \tag{36.7-9c}$$

$$A_{n+1} = A_n + Y_{n+1}\,(I_{n+1} + A_n) \quad \to W(y) \tag{36.7-9d}$$

Implementation in pari/gp:

```
1   w2(y, N=7)=
2   {
3       local(in, y2, A);
4       /* correct to order = 2^N-1 */
5       in = 1/2; /* 1/2^k * (1+y+y^2+y^3+...+y^(2^k-1)) */
6       A = y/2;
7       for(k=2, N,
8           in *= (1+y)/2;
9           y *= y;
10          A += y*(in + A);
11      );
12      return( A );
13  }
```

In the Taylor series

$$\begin{aligned} W(y) = {} & 0 + \frac{1}{2}y + \frac{1}{4}y^2 + \frac{3}{4}y^3 + \frac{1}{8}y^4 + \frac{5}{8}y^5 + \frac{3}{8}y^6 + \frac{7}{8}y^7 + \\ & + \frac{1}{16}y^8 + \frac{9}{16}y^9 + \frac{5}{16}y^{10} + \frac{13}{16}y^{11} + \frac{3}{16}y^{12} + \frac{11}{16}y^{13} + \frac{7}{16}y^{14} + \frac{15}{16}y^{15} + \\ & + \frac{1}{32}y^{16} + \frac{17}{32}y^{17} + \frac{9}{32}y^{18} + \frac{25}{32}y^{19} + \frac{5}{32}y^{20} + \frac{21}{32}y^{21} + \frac{13}{32}y^{22} + \frac{29}{32}y^{23} + \\ & + \frac{3}{32}y^{24} + \frac{19}{32}y^{25} + \frac{11}{32}y^{26} + \frac{27}{32}y^{27} + \frac{7}{32}y^{28} + \frac{23}{32}y^{29} + \frac{15}{32}y^{30} + \frac{31}{32}y^{31} + \frac{1}{64}y^{32} + \ldots \end{aligned} \tag{36.7-10}$$

the coefficient of the y^n is the weighted sum of digits $w(n) = \sum_{i=0}^{\infty} 2^{-(i+1)}\, b_i$ where $b_0, b_1, \ldots$ is the base-2 representation of n. Note that the numerator in the n-th coefficient is the reversed binary expansion of n.

The corresponding *weighted sum-of-digits constant* or *revbin constant* is $W := W(1/2)$. Then

$$\begin{aligned} W &= 0.448526550676272378923687721254526097616278813538448133 6\ldots \\ [\text{base } 2] &= 0.011100101101001010100010110100101000101011010010101000 1\ldots \\ [\text{CF}] &= [0, 2, 4, 2, 1, 4, 18, 1, 2, 6, 5, 17, 2, 14, 1, 1, 1, 2, 1, 1, 2, 1, 3, 1, 29, 4, 1, \ldots] \end{aligned} \tag{36.7-11}$$

For the function W we have the following functional relation:

$$W\left(\frac{1}{q}\right) = \frac{1}{2}\left[\frac{1+q}{q}\, W\left(\frac{1}{q^2}\right) + \frac{q}{q^2-1}\right] \tag{36.7-12}$$

36.8 Iterations related to the binary Gray code

36.8.1 Series where coefficients are the Gray code of exponents

We construct a function with Taylor series coefficients that are the binary Gray code of the exponent of y. A list of the Gray codes is given below (see section 1.16 on page 40):

```
k ==
0 1 2 3 4 5 6 7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
0 1 3 2 6 7 5 4 12 13 15 14 10 11  9  8 24 25 27 26 30 31 29 28 20 21 23 22 18 19 17 16
== graycode(k)
```

The sequence of Gray codes is entry A003188 of [240]. Define the function $G(y)$ as the limit of the iteration

$$\begin{aligned}
F_1 &= y, \qquad B_1 = 1, \qquad Y_1 = y, \qquad I_1 = 1+y && \text{(36.8-1a)}\\
Y_n &= Y_{n-1}^2 && \text{(36.8-1b)}\\
F_n &= (F_{n-1} \qquad\qquad) + Y_n\,(B_{n-1} + 2\,I_{n-1}) \quad \to G(y) && \text{(36.8-1c)}\\
B_n &= (F_{n-1} + 2\,I_{n-1}) + Y_n\,(B_{n-1} \qquad\qquad) && \text{(36.8-1d)}\\
I_n &= I_{n-1}\,(1+Y_n) && \text{(36.8-1e)}
\end{aligned}$$

Implementation in pari/gp:

```
1   gg(y, N=15)=
2   {
3       local(t, ii, F, B, Fp, Bp);
4       /* correct up to order 2^N-1 */
5       F=0+y;  B=1+0;  ii=1+y;
6       for(k=2,N,
7         y *= y;
8         ii *= 2;  /* delete for sum of digits */
9         Fp = (F     ) +  y * (B + ii);
10        Bp = (F + ii) +  y * (B     );
11        F = Fp;  B = Bp;
12        ii *= (1+y);
13      );
14      return( F )
15  }
```

In the algorithm `F` contains the approximation so far and `B` contains the reversed polynomial:

```
 ---- k = 1 :
F = (y)
B = (1)
 ---- k = 2 :
F = (2y^3+3y^2+y)
B = (y^2+3y+2)
 ---- k = 3 :
F = (4y^7+5y^6+7y^5+6y^4+2y^3+3y^2+y)
B = (y^6+3y^5+2y^4+6y^3+7y^2+5y+4)
 ---- k = 4 :
F = (8y^15+9y^14+11y^13+10y^12+14y^11+15y^10+13y^9+12y^8+4y^7+5y^6+7y^5+6y^4+2y^3+3y^2+y)
B = (y^14+3y^13+2y^12+6y^11+7y^10+5y^9+4y^8+12y^7+13y^6+15y^5+14y^4+10y^3+11y^2+9y+8)
```

We obtain the series

$$\begin{aligned}
G(y) = {} & 0 + 1y + 3y^2 + 2y^3 + 6y^4 + 7y^5 + 5y^6 + 4y^7 + \\
& +12y^8 + 13y^9 + 15y^{10} + 14y^{11} + 10y^{12} + 11y^{13} + 9y^{14} + 8y^{15} + \\
& +24y^{16} + 25y^{17} + 27y^{18} + 26y^{19} + 30y^{20} + 31y^{21} + 29y^{22} + 28y^{23} + \\
& +20y^{24} + 21y^{25} + 23y^{26} + 22y^{27} + 18y^{28} + 19y^{29} + 17y^{30} + 16y^{31} + \ldots
\end{aligned} \qquad \text{(36.8-2)}$$

We define the *Gray code constant* as $G := G(1/2)$:

$$\begin{aligned}
G &= 2.302218287787689301229333006391310761000431077704369505\ldots && \text{(36.8-3)}\\
[\text{base } 2] &= 10.010011010101111000101101011111100100110100011110001 01\ldots\\
[\text{CF}] &= [2, 3, 3, 4, 4, 1, 4, 4, 1, 2, 1, 1, 1, 2, 24, 205, 1, 4, 2, 2, 1, 1, 4, 10, 8, 1, 9, 1, \ldots]
\end{aligned}$$

For the function G we have

$$G\left(\frac{1}{q^2}\right) = \frac{1}{4}\left[\frac{q}{q+1}\,G\left(\frac{1}{q}\right) + \frac{q}{q-1}\,G\left(-\frac{1}{q}\right)\right] \tag{36.8-4a}$$

$$G\left(\frac{1}{q}\right) = \frac{2\,(q+1)}{q}\,G\left(\frac{1}{q^2}\right) + \frac{q^2}{(q-1)\,(q^2+1)} \tag{36.8-4b}$$

$$G(y) = 2\,(1+y)\,G(y^2) + \frac{y}{(1-y)\,(1+y^2)} \tag{36.8-4c}$$

$$G\left(-\frac{1}{q}\right) = \frac{2\,(q-1)}{q}\,G\left(\frac{1}{q^2}\right) - \frac{q^2}{(q+1)\,(q^2+1)} \tag{36.8-4d}$$

$$G\left(\frac{1}{q}\right) = \frac{2\,(q^2+1)}{q\,(q-1)}\,G\left(-\frac{1}{q^2}\right) + \frac{q^2\,(q^4+4\,q^3+4\,q+1)}{(q^4+1)\,(q^2+1)\,(q-1)} \tag{36.8-4e}$$

$$G\left(\frac{1}{q}\right) = \frac{1}{2\,q^2}\left[\frac{(q+1)\,(q^4+4\,q^3+4\,q+1)}{(q^2+1)}\,G\left(\frac{1}{q^2}\right) - \frac{(q^4+1)}{(q-1)}\,G\left(-\frac{1}{q^2}\right)\right] \tag{36.8-4f}$$

The function $G(y)$ can be expressed as

$$G(y) = \frac{1}{1-y}\sum_{k=0}^{\infty}\frac{2^k\,y^{2^k}}{1+y^{2^{k+1}}} \tag{36.8-5}$$

36.8.2 Differences of the Gray code

If one defines $F(y) = (1-y)\,G(y)$ to obtain the successive differences of the Gray code itself, the Taylor coefficients are powers of two in magnitude:

$$F(y) = 0 + y + 2y^2 - y^3 + 4y^4 + y^5 - 2y^6 - y^7 + 8y^8 + y^9 + 2y^{10} - y^{11} - 4y^{12} + y^{13} \pm \ldots \tag{36.8-6}$$

We have

$$F\,(y) = 2\,F\left(y^2\right) + \frac{y}{1+y^2} \tag{36.8-7}$$

Now, as $y/(1+y) = q/(1+q)$ for $q = 1/y$,

$$F\,(y) = F\left(\frac{1}{y}\right) = \sum_{k=0}^{\infty}2^k\,\frac{y^{2^k}}{1+y^{2^{k+1}}} \tag{36.8-8}$$

Thereby $F(y)$ can be computed everywhere except on the unit circle. The sum

$$\sum_{k=0}^{\infty}2^k\,\frac{y^{2^k}}{1-y^{2^{k+1}}} \tag{36.8-9}$$

leads to a series with coefficients

```
0 1 2 1 4 1 2 1 8 1 2 1 4 1 2 1 16 1 2 1 4 1 2 1 8 1 2 1 4 1 2 1 32 1 ...
```

corresponding to the (exponential) version of the ruler function which is defined as the highest power of two that divides n. The ruler function (see section 36.4 on page 720)

```
. 0 1 0 2 0 1 0 3 0 1 0 2 0 1 0 4 0 1 0 2 0 1 0 3 0 1 0 2 0 1 0 5 0 1 ...
```

is the base-2 logarithm of that series.

36.8.3 Sum of Gray code digits

The sequence of the sum of digits of the Gray code of $k \geq 0$ is (entry A005811 in [240]):

```
0 1 2 1 2 3 2 1 2 3 4 3 2 3 2 1 2 3 4 3 4 5 4 3 2 3 4 3 2 3 2 1 2 3 ...
```

Omit the factor 2 in relations 36.8-1c and 36.8-1d on page 729. That is, in the implementation simply remove the line

```
ii *= 2;  /* delete for sum of digits */
```

Let $R(y)$ be the corresponding function, and define the *sum of Gray code digits constant* as $R := R(1/2)/2$, then

$$\begin{aligned} R &= 0.701472376403734520735595521064133208822798986165421295 4\ldots \quad (36.8\text{-}10)\\ [\text{base } 2] &= 0.101100111001001110110001100100111011001110010001101100 0\ldots \\ [\text{CF}] &= [0, 1, 2, 2, 1, 6, 10, 1, 9, 53, 1, 1, 3, 10, 1, 2, 1, 3, 2, 14, 2, 1, 2, 1, 3, 4, 2, \\ &\quad 1, 34, 1, 1, 3, 1, 1, 109, 1, 1, 4, 2, 9, 1, 642, 51, 4, 3, 2, 2, 2, 2, 1, 2, 3, \ldots] \end{aligned}$$

One finds:

$$R\left(\frac{1}{q^2}\right) = \frac{1}{2}\left[\frac{q}{q+1}\,R\left(\frac{1}{q}\right) + \frac{q}{q-1}\,R\left(-\frac{1}{q}\right)\right] \quad (36.8\text{-}11a)$$

$$R\left(\frac{1}{q}\right) = \frac{q+1}{q}\,R\left(\frac{1}{q^2}\right) + \frac{q^2}{q^3-q^2+q-1} \quad (36.8\text{-}11b)$$

$$R\left(-\frac{1}{q}\right) = \frac{q-1}{q}\,R\left(\frac{1}{q^2}\right) - \frac{q^2}{q^3+q^2+q+1} \quad (36.8\text{-}11c)$$

The function $R(y)$ can be expressed as

$$R(y) = \frac{1}{1-y}\sum_{k=0}^{\infty}\frac{y^{2^k}}{1+y^{2^{k+1}}} \quad (36.8\text{-}12)$$

Define the constant P as $P = (R+1)/2$

$$\begin{aligned} P &= 0.850736188201867260367797760532066604411399493082710647 7\ldots \quad (36.8\text{-}13)\\ [\text{base } 2] &= 0.110110011100100111011000110010011101100111001000110110 0\ldots \\ [\text{CF}] &= [0, 1, 5, 1, 2, 3, 21, 1, 4, 107, 7, 5, 2, 1, 2, 1, 1, 2, 1, 6, 1, 2, 6, 1, 1, 8, 1, \\ &\quad 2, 17, 3, 1, 1, 3, 1, 54, 3, 1, 1, 1, 2, 1, 4, 2, 321, 102, 2, 6, 1, 4, 1, 5, 2, \ldots] \end{aligned}$$

Its binary expansion is the *paper-folding sequence* (or *dragon curve sequence*), entry A014577 of [240].

36.8.4 Differences of the sum of Gray code digits

Now define $E(y) = (1-y)\,R(y)$ to obtain the differences of the sum of Gray code digits. From this definition (and relation 36.8-12) one sees that

$$E(y) = \sum_{k=0}^{\infty}\frac{y^{2^k}}{1+y^{2^{k+1}}} \quad (36.8\text{-}14)$$

All Taylor coefficients (except for the constant term) are either plus or minus one:

$$E(y) = 0 + y + y^2 - y^3 + y^4 + y^5 - y^6 - y^7 + y^8 + y^9 + y^{10} - y^{11} - y^{12} + y^{13} \pm \ldots \quad (36.8\text{-}15)$$

We have

$$E\left(\frac{1}{q}\right) = E\left(\frac{1}{q^2}\right) + \frac{q}{q^2+1} \tag{36.8-16a}$$

$$E(y) = E\left(y^2\right) + \frac{y}{y^2+1} \tag{36.8-16b}$$

(use $\frac{y}{y^2+1} = \frac{q}{q^2+1}$ where $q = \frac{1}{y}$ for the latter relation), thereby

$$E(y) = E\left(\frac{1}{y}\right) \tag{36.8-17}$$

So we can compute $E(y)$ everywhere except on the unit circle. For $y < 1$ compute $E(y)$ by the iteration

$$L_0 = 0, \qquad R_0 = 1, \qquad Y_0 = y \tag{36.8-18a}$$

$$L_{n+1} = L_n + Y_n R_n \quad \to E(y) \tag{36.8-18b}$$

$$R_{n+1} = (L_n + 1) + Y_n (R_n - 2) \tag{36.8-18c}$$

$$Y_{n+1} = Y_n^2 \tag{36.8-18d}$$

Implementation in pari/gp:

```
1   ge(y, N=7)=
2   {
3       local(L, R, Lp, Rp);
4       /* correct up to order 2^N-1 */
5       L=0;  R=1;
6       for(k=2, N,
7         Lp = (L    ) +  y * (R    );
8         Rp = (L + 1) +  y * (R - 2);
9         L = Lp;  R = Rp;
10        y *= y;
11      );
12      return( L + y*R )
13  }
```

The symbolic representations of the polynomials `L` and `R` should make the underlying idea clear:

```
L = 0
R = +

L = 0+
R = +-

L = 0++-
R = ++--

L = 0++-++--
R = +++--+--

L = 0++-++--+++--+--
R = +++-++---++--+--

L = 0++-++--+++--+--+++-++---++--+--
R = +++-++--+++--+---++-++---++--+--

L = 0++-++--+++--+--+++-++---++--+--+++-++--+++--+---++-++---++--+--
R = +++-++--+++--+--+++-++---++--+---++-++--+++--+---++-++---++--+--
```

The limit of the sequence `L` is entry A034947 of [240]. It is a signed version of the *paper-folding sequence*. The sequence after the initial zero is identical to the sequence of the Kronecker symbols $\left(\frac{-1}{n}\right)$ for $n = 0, 1, 2, 3, \ldots$, see section 37.8 on page 771. Quick verification:

```
? for(n=1,88,print1(if(-1==kronecker(-1,n),"-","+")))
  ++-++--+++--+--+++-++---++--+--+++-++--+++--+---++-++---++--+--+++-++---++--+--+++-++---
```

We note that the (divisionless) algorithm is fast in practice, we compute the constant R as $R = E(1/2)$ to more than 600 thousand decimal digits:

```
? N=21 ; \\ number of iterations
? B=2^N-1  \\ precision in bits
  2097151
? D=ceil(B * log(2)/log(10)) + 1  \\ precision in decimal digits
  631307
? default(realprecision,D);
```

```
? r=ge(1.0/2, N)
  0.701472376403734520735595521064 [... + > 600k digits]
? ##
  ***   last result computed in 2,726 ms.
? r-ge(1.0/2, N+2)  \\ checking precison
  -3.744192708419695540526787319381583798 4 E-631306
? p=(1+r)/2
  0.850736188201867260367797760532 [... + > 600k digits]
```

36.8.5 Weighted sum of Gray code digits

Define $H(y)$ by its Taylor series $H(y) = \sum_{k=1}^{\infty} h(k)y^k$ where $h(k)$ is the weighted sum of digits of the Gray code:

```
1   wgs(k)=
2   {
3       local(g,t,s);  s=0;  g=gray(k);
4       for(n=0,33, s+=if(bittest(g,n),1/2^(n+1),0));
5       return(s);
6   }
```

```
? for(k=1,33,  print1(" ",wgs(k)))
  1/2
  3/4 1/4
  3/8 7/8 5/8 1/8
 3/16 11/16 15/16 7/16 5/16 13/16 9/16 1/16
 3/32 19/32 27/32 11/32 15/32 31/32 23/32 7/32 5/32 21/32 29/32 13/32 9/32 25/32 17/32 1/32
 3/64 35/64 ...
```

An iteration for the computation of $H(y)$ is:

$$\begin{aligned}
F_1 &= y/2, \qquad B_1 = 1/2, \qquad Y_1 = y, \qquad I_1 = (1+y)/2 && \text{(36.8-19a)}\\
Y_n &= Y_{n-1}^2 && \text{(36.8-19b)}\\
K_n &= I_{n-1}/2 && \text{(36.8-19c)}\\
F_n &= (F_{n-1} \qquad) + Y_n\,(B_{n-1} + K_n) \quad \to H(y) && \text{(36.8-19d)}\\
B_n &= (F_{n-1} + K_n) + Y_n\,(B_{n-1} \qquad) && \text{(36.8-19e)}\\
I_n &= K_n\,(1 + Y_n) && \text{(36.8-19f)}
\end{aligned}$$

Implementation in pari/gp:

```
1   gw(y, N=11)=
2   {
3       local(t, ii, F, B, Fp, Bp);
4       /* correct up to order 2^N-1 */
5       F=0+y;  B=1+0;  ii=1+y;
6       ii /= 2;  F /= 2;  B /= 2;
7       for(k=2,N,
8         y *= y;
9         ii /= 2;
10        Fp = (F     ) +  y * (B + ii);
11        Bp = (F + ii) +  y * (B     );
12        F = Fp;  B = Bp;
13        ii *= (1+y);
14      );
15      return( F )
16  }
```

We define the *weighted sum of Gray code digits constant* as $H := H(1/2)$, then

$$\begin{aligned}
H &= 0.533700488639284991958880481482124285854919322545611891 1\ldots && \text{(36.8-20)}\\
[\text{base } 2] &= 0.100010001010000010011000011000001001000010100000011110 0\ldots\\
[\text{CF}] &= [0, 1, 1, 6, 1, 11, 4, 5, 6, 1, 13, 1, 3, 1, 18, 5, 77, 1, 2, 2, 3, 1, 2, 1, 1, \ldots]
\end{aligned}$$

We have:

$$H\left(\frac{1}{q^2}\right) = \frac{q}{q+1}\,H\left(\frac{1}{q}\right) + \frac{q}{q-1}\,H\left(-\frac{1}{q}\right) \tag{36.8-21a}$$

$$H\left(\frac{1}{q}\right) = \frac{1}{2}\left[\frac{q+1}{q}\,H\left(\frac{1}{q^2}\right) + \frac{q^2}{q^3-q^2+q-1}\right] \tag{36.8-21b}$$

$$H\left(-\frac{1}{q}\right) = \frac{1}{2}\left[\frac{q-1}{q}\,H\left(\frac{1}{q^2}\right) - \frac{q^2}{q^3+q^2+q+1}\right] \tag{36.8-21c}$$

36.9 A function that encodes the Hilbert curve

We define a function $H(y)$ by the following iteration:

$$H_1 = +i\,y + 1\,y^2 - i\,y^3 \tag{36.9-1a}$$

$$R_1 = +i\,y - 1\,y^2 - i\,y^3 \tag{36.9-1b}$$

$$Y_1 = y \tag{36.9-1c}$$

$$Y_{n+1} = Y_n^4 \tag{36.9-1d}$$

$$H_{n+1} = -i\,R_n + Y_n\,(+i + H_n + Y_n\,(+1 + H_n + Y_n\,(-i + i\,R_n))) \tag{36.9-1e}$$

$$R_{n+1} = +i\,H_n + Y_n\,(+i + R_n + Y_n\,(-1 + R_n + Y_n\,(-i - i\,H_n))) \tag{36.9-1f}$$

$$H_{2\,n} \to H(y) \tag{36.9-1g}$$

As the real and imaginary parts are swapped with each step we agree on iterating an even number of times. The resulting function $H(y)$ is

$$H(y) = 0 + y + i\,y^2 - y^3 + i\,y^4 + i\,y^5 + y^6 - i\,y^7 + y^8 + i\,y^9 + y^{10} - i\,y^{11} \pm \ldots \tag{36.9-2}$$

The coefficients of the series are ± 1 and $\pm i$, except for the constant term which is zero. If the sequence of coefficients is interpreted as follows:

```
 0  :=  goto start  '0'
+1  :=  move right  '>'
-1  :=  move left   '<'
+i  :=  move up     '^'
-i  :=  move down   'v'
```

Then, symbolically:

```
H = 0>^<^^>v>^>vv<v>>^>v>>^<^>^<<v<^^^>v>>^<^>^<<v<^<<v>vv<^<v<^^>^< ...
```

Follow the signs to walk along the Hilbert curve, see figure 1.20-A on page 56. An implementation is

```
1   hh(y, N=4)=
2   {
3       /* correct to order = 4^N-1 */
4       local(H, R, tH, tR);
5       H= +I*y + 1*y^2 -I*y^3;  R= +I*y - 1*y^2 -I*y^3;
6       for(k=2, N,
7           y=y^4;
8           tH = -I*R + y*(+I + H + y*(+1 + H + y*(-I + I*R)));
9           tR = +I*H + y*(+I + R + y*(-1 + R + y*(-I - I*H)));
10          H=tH;  R=tR;
11      );
12      return( H );
13  }
```

The value of $H(y)$ for $y < 1$ gives the limiting point in the complex plane when the walk according to the coefficients is done with decreasing step lengths: step number k has step length y^k. The least positive y where the real and imaginary part of the endpoint are equal is $y_1 = 0.5436890126920\ldots$. It turns out that y_1 is the real solution of the polynomial $y^3 + y^2 + y - 1$. On might suspect that $M(y) := \mathfrak{Re}\,H(y) - \mathfrak{Im}\,H(y)$, the difference between the real and the imaginary part of $H(y)$, has the

factor y^3+y^2+y-1. Indeed we have $M(y) = y\,(y^3+y^2+y-1)\,(y^{12}+y^8+y^4-1)\cdots$, and a similar statement is true for the $P(y) := \mathfrak{Re}\,H(y) + \mathfrak{Im}\,H(y)$. We use this observation for the construction of a simplified and quite elegant algorithm for the computation of $H(y)$.

36.9.1 A simplified algorithm

Define the function $P(y)$ as the result of the iteration

$$\begin{aligned} Y_1 &= y, \qquad P_1 = 1 && \text{(36.9-3a)}\\ P_{n+1} &= P_n\,\left(+1+Y_n+Y_n^2-Y_n^3\right) && \text{(36.9-3b)}\\ Y_{n+1} &= Y_n^4 && \text{(36.9-3c)}\\ P_n - 1 &\to P(y) && \text{(36.9-3d)} \end{aligned}$$

and the function $M(y)$ by

$$\begin{aligned} Y_1 &= y, \qquad M_1 = 1 && \text{(36.9-4a)}\\ M_{n+1} &= M_n\,\left(-1+Y_n+Y_n^2+Y_n^3\right) && \text{(36.9-4b)}\\ Y_{n+1} &= Y_n^4 && \text{(36.9-4c)}\\ y\,M_n &\to M(y) && \text{(36.9-4d)} \end{aligned}$$

Now the function $H(y)$ can be computed as

$$H(y) = \frac{1}{2}\left[(P(y)+M(y)) + i\,(P(y)-M(y))\right] \qquad \text{(36.9-5)}$$

The following implementations compute the series up to order 4^N-1:

```
fpp(y, N=4)=
{
    local( t, Y );
    t = 1;  Y=y;
    for (k=1, N,  t *= (+1+Y+Y^2-Y^3);  Y=Y^4; );
    return( t-1 );
}

fmm(y, N=4)=
{
    local( t, Y );
    t = 1;  Y=y;
    for (k=1, N,  t*= (-1+Y+Y^2+Y^3);  Y=Y^4; );
    return( t*y-Y );
}

hhpm(y, N=4)=
{
    local( tp, tm );
    tp = fpp(y);
    tm = fmm(y);
    return( ((tp+tm) + I*(tp-tm))/2 );
}
```

With a routine `tdir()` that prints a power series with coefficients $\in \{-1,\,0,\,+1\}$ symbolically we obtain:

```
? N=4;
? tdir(fpp(y));tdir(fmm(y));
  0++-+++-++----++++-+++-+++----++++-+++-++----+---+---+---++++-++
  0+----+++-+++-+++-++++---+---+----++++---+---+----++++---+---+----

? tdir((fpp(y)+fmm(y))/2);tdir((fpp(y)-fmm(y))/2);
  0+0-00+0+0+00-0++0+0++0-0+0--0-000+0++0-0+0--0-0--0+00-0-0-00+0-00
  00+0++0-0+0--0-00+0-00+0+0+00-0+++0-00+0+0+00-0+00-0--0+0-0++0+0++
```

The n-th coefficient of the Taylor series of $P(y)$ equals the parity of the number of threes in the radix-4 representation of n. This fact can be used for an efficient bit level algorithm, see section 1.20 on page 56.

The coefficients of the power series of the functions P and M can be computed with a string substitution engine, see figure 36.9-A.

```
Number of symbols = 2
Start: +
Rules:
  + --> +++-
  - --> ---+
-------------
0:   (#=1)
  +
1:   (#=4)
  +++-
2:   (#=16)
  +++-+++-+++----+
3:   (#=64)
  +++-+++-+++----++++-+++-+++----++++-+++-+++----+---+---+---++++-
4:   (#=256)
  +++-+++-+++----++++-+++-+++----++++-+++-+++----+---+---+---++++-+++-+++-++ ...

  0++-+++-+++----++++-+++-+++----++++-+++-+++----+---+---+---++++-++ == fpp()
```

```
Number of symbols = 2
Start: +
Rules:
  + --> -+++
  - --> +---
-------------
0:   (#=1)
  +
1:   (#=4)
  -+++
2:   (#=16)
  +----+++-+++-+++
3:   (#=64)
  -++++----+---+---+----+++-+++-++++----+++-+++-++++----+++-+++-+++
4:   (#=256)
  +----+++-+++-+++-++++---+---+----++++---+---+----++++---+---+----++++---+- ...

 0+----+++-+++-+++-++++---+---+----++++---+---+----++++---+---+---- == fmm()
```

Figure 36.9-A: Computation of the power series of the functions P (top) and M (bottom) with a string substitution engine.

36.9.2 The turns of the Hilbert curve

We compute a function with series coefficients $\in \{-1, 0, +1\}$ that correspond to the turns of the Hilbert curve. We use $+1$ for a right turn, -1 for a left turn and zero for no turn. The sequences of turns starts as

```
0--+0++--++0+--0-++-0--++--0-++00++-0--++--0-++-0--+0++--++0+--+ \
0++-0--++--0-++0+--+0++--++0+--00--+0++--++0+--+0++-0--++--0-++- \
-++-0--++--0-++0+--+0++--++0+--00--+0++--++0+--+0++-0--++--0-++0 \
+--+0++--++0+--0-++-0--++--0-++00++-0--++--0-++-0--+0++--++0+-- ...
```

The computation is slightly tricky:

```
1   hht(y, N=4)=
2   {
3       /* correct to order = 4^N-1 */
4       local( t, Y, F, s, p );
5       t = 1;  Y=y;  p = 1;
6       F = y + y^2;
7       for(k=2, N,
8           Y = Y^4;
9           t = -F + Y*F + Y^2*F  - Y^3*F;
10          p *= 4;
11          if ( 0==(k%2),
12              t +=  y^(1*p-1);
13              t +=  y^(3*p);
14              t -=  (y+1)*y^(2*p-1);
15              , /* else */
16              t +=  y^(1*p);
17              t +=  y^(3*p-1);
18          );
19          F = t;
20      );
21      if ( 1==N%2,  F = -F );  \\ same result for even and odd N
22      return( F );
23  }
```

36.10 Sparse variants of the inverse

36.10.1 A fourth order iteration

Define the function $F(y)$ as the result of the iteration

$$\begin{aligned} F_0 &= 1, \qquad Y_0 = y && \text{(36.10-1a)}\\ F_{n+1} &= F_n\,(1+Y_n) \quad \to F(y) && \text{(36.10-1b)}\\ Y_{n+1} &= Y_n^4 && \text{(36.10-1c)} \end{aligned}$$

The terms of the continued fraction of $F(1/10)$ grow doubly exponentially:

```
[1, 9, 1, 90, 110, 909000, 11001100000, 9090000090900000000000,
 1100110000000000110011000000000000000000000000,
 909000009090000000000000000000000090900000909000000000000000000000000000000000000000000000000000000
  ... ]
```

Let $[t_0, t_1, t_2, \ldots]$ be the continued fraction of $F(1/q)$, then

$$\begin{aligned} t_0 &= 1 && \text{(36.10-2a)}\\ t_1 &= q-1 && \text{(36.10-2b)}\\ t_2 &= 1 && \text{(36.10-2c)}\\ t_3 &= q^2-q && \text{(36.10-2d)}\\ t_4 &= q^2+q && \text{(36.10-2e)}\\ t_5 &= q^6-q^5+q^4-q^3 && \text{(36.10-2f)}\\ t_6 &= q^{10}+q^9+q^6+q^5 && \text{(36.10-2g)}\\ t_7 &= q^{22}-q^{21}+q^{20}-q^{19}+q^{14}-q^{13}+q^{12}-q^{11} && \text{(36.10-2h)}\\ t_8 &= q^{42}+q^{41}+q^{38}+q^{37}+q^{26}+q^{25}+q^{22}+q^{21} && \text{(36.10-2i)}\\ t_9 &= q^{86}-q^{85}+q^{84}-q^{83}+q^{78}-q^{77}+q^{76}-q^{75}+ && \text{(36.10-2j)}\\ &\quad +q^{54}-q^{53}+q^{52}-q^{51}+q^{46}-q^{45}+q^{44}-q^{43} \end{aligned}$$

For $j \geq 4$ we have

$$\frac{t_j}{t_{j-2}} = q^{2J}+q^J = q^J\,(q^J+1) \quad \text{where} \quad J = 2^{j-4} \qquad \text{(36.10-3)}$$

A functional relation for F is given by

$$F(y)\,F(y^2) = \frac{1}{1-y} \qquad \text{(36.10-4)}$$

The relation is (mutatis mutandis) also true for the truncated product. The binary expansions of $F(y)$, $F(y^2)$ and their product for $y = 1/2$ are (dots for zeros):

```
11..11.........11..11.....................................................
1.1.....1.1.....................1.1.....1.1.....................................
1111111111111111111111111111111111111111111111111111111111111111111111....
```

Since the expansions are palindromes also their correlation is a sequence of ones. Now set $F_k := F(y^k)$ then, by relations 36.10-4 and 36.1-16e on page 716 (for $k = 2$, 3, and 5),

$$\begin{aligned} 0 &= F_1^2\,F_2 - 2\,F_1\,F_2\,F_4 + F_4 && \text{(36.10-5a)}\\ 0 &= F_1^3\,F_2^3 - F_3\,F_6\,\left(3\,F_1^2\,F_2^2 - 3\,F_1\,F_2 + 1\right) && \text{(36.10-5b)}\\ 0 &= F_1^5\,F_2^5 - F_{10}\,F_5\,\left(5\,F_1^4\,F_2^4 - 10\,F_1^3\,F_2^3 + 10\,F_1^2\,F_2^2 - 5\,F_1\,F_2 + 1\right) && \text{(36.10-5c)} \end{aligned}$$

We note that for power series over GF(2) relation 36.10-4 becomes $F(y)^3 = 1/(1-y)$. That is, $F(y) = \sqrt[3]{1-y}$. In general, an iteration for the inverse (2^w-1)-st root is obtained by replacing relation 36.10-1c with $Y_{k+1} = Y_k^{2^w}$.

36.10.2 A different fourth order iteration

To define the function $F(y)$ we modify the third order iteration for $\frac{1}{1+y}$

```
1   inv3m(y, N=6)=  /* third order --> 1/(1+y) */
2   { /* correct to order 3^N */
3       local(T);
4       T = 1;
5       for(k=1, N,
6         T *= ( 1 - y + y^2 );
7         y = y^3;
8       );
9       return( T );
10  }
```

to obtain a fourth-order iteration:

```
1   f43(y, N=6)=
2   { /* correct to order 4^N */
3       local(T, yt);
4       T = 1;
5       for(k=1, N,
6         T *= ( 1 - y + y^2 );
7         y = y^4;  /* ! */
8       );
9       return( T );
10  }
```

That is,

$$\begin{aligned} F_0 &= 1, \qquad Y_0 = y && \text{(36.10-6a)}\\ F_{n+1} &= F_n\,(1 - Y_n + Y_n^2) \quad \to F(y) && \text{(36.10-6b)}\\ Y_{n+1} &= Y_n^4 && \text{(36.10-6c)} \end{aligned}$$

The first few terms of the power series are

$$F(y) = 1 - y + y^2 - y^4 + y^5 - y^6 + y^8 - y^9 + y^{10} - y^{16} + y^{17} - y^{18} + y^{20} - y^{21} \pm \ldots \quad \text{(36.10-7)}$$

Let $[t_0, t_1, t_2, \ldots]$ be the continued fraction of $F(1/q)$, then

$$\begin{aligned} t_0 &= 0 && \text{(36.10-8a)}\\ t_1 &= 1 && \text{(36.10-8b)}\\ t_2 &= q && \text{(36.10-8c)}\\ t_3 &= q && \text{(36.10-8d)}\\ t_4 &= q^2 - q && \text{(36.10-8e)}\\ t_5 &= q^4 + q^3 - q - 1 && \text{(36.10-8f)}\\ t_6 &= q^8 - q^7 + q^6 - q^4 + q^3 - q^2 && \text{(36.10-8g)}\\ t_7 &= q^{16} + q^{15} - q^{13} + q^{11} + q^{10} - q^8 - q^7 + q^5 - q^3 - q^2 && \text{(36.10-8h)} \end{aligned}$$

For $j \geq 6$ we have

$$\frac{t_j}{t_{j-2}} = q^{6J} + q^{4J} + q^{3J} + q^J = q^J\,(q^J + 1)\,(q^{2J} - q^J + 1)\,(q^{2J} + 1) \quad \text{(36.10-9)}$$

where $J = 2^{j-6}$. The terms of the continued fraction of $F(1/q)$ for integer q grow doubly exponentially:

```
? contfrac(f43(0.5))
  [0, 1, 2, 2, 2, 21, 180, 92820, 3032435520, 26126907554432455680,
   240254294248527099500117907463345274880,
   164001256750215347067944129734442019102853751066678216639025390799096507269120,
   ... ]
/* number of decimal digits of the terms in the CF: */
  [-, 1, 1, 1, 1, 2, 3, 5, 10, 20, 39, 78, 154, 309, ... ]
```

By construction,

$$F(y) = \left(1 - y + y^2\right)\, F\left(y^4\right) \tag{36.10-10a}$$
$$F(-y) = \left(1 + y + y^2\right)\, F\left(y^4\right) \tag{36.10-10b}$$

The equivalent forms with $y = 1/q$ are

$$F\left(\frac{1}{q}\right) = \frac{q^2 - q + 1}{q^2}\, F\left(\frac{1}{q^4}\right) \tag{36.10-11a}$$
$$F\left(-\frac{1}{q}\right) = \frac{q^2 + q + 1}{q^2}\, F\left(\frac{1}{q^4}\right) \tag{36.10-11b}$$

Now $q^2 - q + 1 = p^2 - p + 1$ if $p = 1 - q$, so

$$F\left(1 - \frac{1}{q}\right) = \frac{q^2 - q + 1}{q^2}\, F\left(\left(1 - \frac{1}{q}\right)^4\right) \quad \text{where} \quad q > 1 \tag{36.10-12a}$$
$$F\left(1 + \frac{1}{q}\right) = \frac{q^2 + q + 1}{q^2}\, F\left(\left(1 + \frac{1}{q}\right)^4\right) \quad \text{where} \quad q < -1 \tag{36.10-12b}$$

Adding relations $\alpha\times$ (36.10-11a) and $\beta\times$ (36.10-12a) and simplifying gives

$$\frac{\alpha\, F(y) + \beta\, F(1-y)}{\alpha\, F(y^4) + \beta\, F((1-y)^4)} = y^2 - y + 1 \quad \text{where} \quad \alpha, \beta \in \mathbb{C} \tag{36.10-13}$$

36.10.3 A sixth order iteration

Define the function $F(y)$ by the iteration

$$F_0 = 1, \qquad Y_0 = y \tag{36.10-14a}$$
$$F_{n+1} = F_n\left(1 + Y_n + Y_n^2\right) \quad \to F(y) \tag{36.10-14b}$$
$$Y_{n+1} = Y_n^6 \tag{36.10-14c}$$

Let $[t_0, t_1, t_2, \ldots]$ be the continued fraction of $F(1/q)$, then

$$t_0 = 1 \tag{36.10-15a}$$
$$t_1 = q - 1 \tag{36.10-15b}$$
$$t_2 = q + 1 \tag{36.10-15c}$$
$$t_3 = q^2 - q \tag{36.10-15d}$$
$$t_4 = q^{10} + q^9 + q^8 + q^4 + q^3 + q^2 \tag{36.10-15e}$$
$$t_5 = q^8 - q^7 + q^5 - q^4 \tag{36.10-15f}$$
$$t_6 = q^{64} + q^{63} + q^{62} + q^{58} + q^{57} + q^{56} + q^{52} + q^{51} + q^{50} + \\ + q^{28} + q^{27} + q^{26} + q^{22} + q^{21} + q^{20} + q^{16} + q^{15} + q^{14} \tag{36.10-15g}$$
$$t_7 = q^{44} - q^{43} + q^{41} - q^{40} + q^{26} - q^{25} + q^{23} - q^{22} \tag{36.10-15h}$$
$$t_8 = q^{388} + q^{387} + \ldots + q^{87} + q^{86} \tag{36.10-15i}$$

For $j \geq 4$ we have

$$\frac{t_j}{t_{j-2}} = \begin{cases} q^J + q^{J/2} & \text{if} \quad j \quad \text{odd} \\ \left(q^{10J} + q^{9J} + q^{8J} + q^{4J} + q^{3J} + q^{2J}\right) / \left(q^J + 1\right) & \text{else} \end{cases} \tag{36.10-16}$$

where $J = 6^{j-4}$.

36.11 An iteration related to the Fibonacci numbers

The so-called *rabbit constant* is

$$\begin{aligned} A &= 0.709803442861291314641787399444575597012502205767860516 9\ldots \quad (36.11\text{-}1)\\ [\text{base } 2] &= 0.101101011011010110101101101101101101101101101101011010110101\ldots \\ [\text{CF}] &= [0, 1, 2, 2, 4, 8, 32, 256, 8192, 2097152, 17179869184, \\ &\quad 36028797018963968, 6189700196426901374495621 12, \ldots] \\ &= [0, 2^0, 2^1, 2^1, 2^2, 2^3, 2^5, 2^8, 2^{13}, \ldots, 2^{F_n}, \ldots] \end{aligned}$$

The sequence of zeros and ones after the decimal point is referred to as *rabbit sequence* or *infinite Fibonacci word*, entry A005614 of [240]. The rabbit sequence can be obtained by starting with a single zero and repeated application of the following substitution rules: simultaneously replace all zeros by one ($0 \to 1$, 'young rabbit gets old') and all ones by one-zero ($1 \to 10$, 'old rabbit gets[1] child'). The evolution is:

```
A0 = 0
A1 = 1
A2 = 10
A3 = 101
A4 = 10110
A5 = 10110101
A6 = 1011010110110
A7 = 10110101101101011010110101
A -->1011010110110101101011010110110101101011010110101...
```

The crucial observation is that each element A_n can be obtained by appending A_{n-2} to A_{n-1}, that is $A_n = A_{n-1}.A_{n-2}$. To compute the value of the rabbit constant in base 2 to N bits precision the whole process requires only copying N bits of data is the minimal conceivable work for a (non-sparse) computation.

We define a function $A(y)$ that has the special value

$$A \;=\; \frac{1}{2}\,A\left(\frac{1}{2}\right) \qquad (36.11\text{-}2)$$

by the equivalent operation for power series. We obtain the iteration

$$\begin{aligned} L_0 &= 0 \qquad R_0 = 1 \qquad l_0 = 1 \qquad r_0 = y && (36.11\text{-}3a)\\ l_{n+1} &= r_n = y^{F_n} && (36.11\text{-}3b)\\ r_{n+1} &= r_n\, l_n = y^{F_{n+1}} && (36.11\text{-}3c)\\ L_{n+1} &= R_n && (36.11\text{-}3d)\\ R_{n+1} &= R_n + r_{n+1}\, L_n = R_n + y^{F_{n+1}}\, L_n \quad \to A(y) && (36.11\text{-}3e) \end{aligned}$$

Here F_n denotes the n-th Fibonacci number (sequence A000045 of [240]):

```
n   0  1  2  3  4  5  6   7   8   9  10  11   12   13   14   15 ...
Fn  0  1  1  2  3  5  8  13  21  34  55  89  144  233  377  610 ...
```

After the n-th step the series in y is correct up to order $F_{n+2} - 1$. That is, the order of convergence equals $\frac{\sqrt{5}+1}{2} \approx 1.6180$. The function $A(y)$ has the power series

$$A(y) \;=\; 1 + y^2 + y^3 + y^5 + y^7 + y^8 + y^{10} + y^{11} + y^{13} + y^{15} + y^{16} + y^{18} + y^{20} + y^{21}\ldots \qquad (36.11\text{-}4)$$

The sequence of exponents of y in the series is entry A022342 of [240], the Fibonacci-even numbers. A pari/gp implementation is

```
1  fa(y, N=10)=
2  {
3      local(t, yl, yr, L, R, Lp, Rp);
4      /* correct up to order fib(N+2)-1 */
5      L=0;  R=1;   yl=1;  yr=y;
```

[1] No sex, no death.

```
 6       for(k=1, N,
 7         t=yr;  yr*=yl;  yl=t;
 8         Lp=R;  Rp=R+yr*L;  L=Lp;  R=Rp;
 9       );
10       return( R )
11   }
```

36.11.1 Fibonacci representation

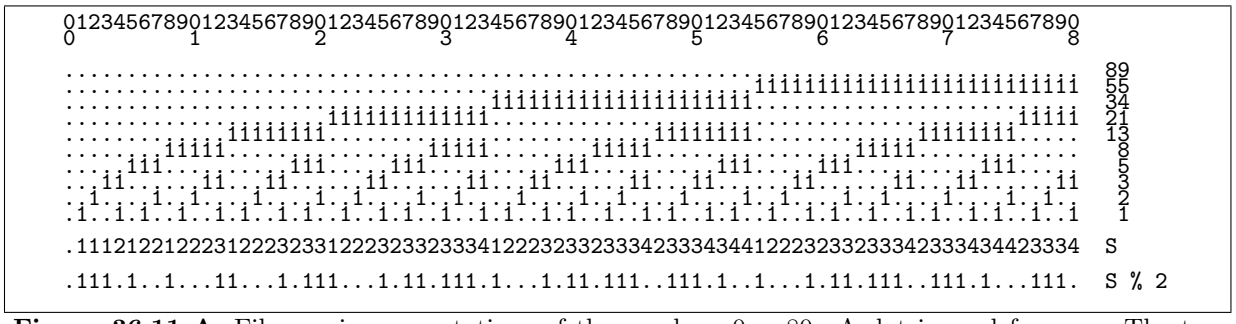

Figure 36.11-A: Fibonacci representations of the numbers $0 \ldots 80$. A dot is used for zero. The two lower lines are the sum of digits and the sum of digits modulo two, the *Fibonacci parity*.

The greedy algorithm to obtain the *Fibonacci representation* (or *Zeckendorf representation*) of an integer repeatedly subtracts the largest Fibonacci number that is greater or equal to it until the number is zero. The Fibonacci representations of the numbers $0 \ldots 80$ are shown in figure 36.11-A.

The sequence of lowest Fibonacci bits is

```
0100101001001010010100100101001001010010100100101001010010100100101001...
```

Interpreted as the binary number $x = 0.1001010010010_2 \ldots = 0.5803931 \ldots$ it turns out that $A = 1 - x/2$ (that is, $x = 2 - A(1/2)$). Alternatively, one can compute the number as $x = A[1, 0, 1, 1/2]$.

The sequence of least significant bits in the Fibonacci representations is entry A003849 of [240]. The sequence of numbers of digits in the Fibonacci representations (second lowest row in figure 36.11-A) is entry A007895 of [240]. This sequence modulo two gives the *Fibonacci parity* (entry A095076). It can be computed by initializing $L_0 = 1$ and changing relation 36.11-3e to

$$R_{n+1} \;=\; R_n - r_{n+1}\, L_n \;=\; R_n - y^{F_{n+1}}\, L_n \quad \to A_p(y) \tag{36.11-5}$$

Let the corresponding function be $A_p(y)$. We define the *Fibonacci parity constant* A_p as

$$\begin{aligned} A_p &= 1 - A_p(1/2)/2 && \text{(36.11-6a)} \\ &= 0.910533470863561763804686886771098007344581229006937645 4\ldots && \text{(36.11-6b)} \\ \text{[base 2]} &= 0.111010010001100010111000101101110100010110111001110100 10\ldots \\ \text{[CF]} &= [0, 1, 10, 5, 1, 1, 1, 3, 4, 2, 6, 25, 4, 5, 1, 1, 3, 5, 1, 3, 2, 1, 1, 1, 3, 1, 3, 22, 1, \\ &\quad 10, 1, 2, 3, 2, 73, 1, 111, 46, 1, 51, 2, 1, 1, 5, 1, 65, 3, 1, 3, 2, 5, 6, 1, 4, 1, 2, \ldots] \end{aligned}$$

The sequence of the Fibonacci representations interpreted as binary numbers is

```
0, 1, 2, 4, 5, 8, 9, 10, 16, 17, 18, 20, 21, 32, 33, 34, 36, 37, 40,
41, 42, 64, 65, 66, 68, 69, 72, 73, 74, 80, 81, 82, 84, 85, 128, 129, ...
```

This is entry A003714 of [240], where the numbers are called *Fibbinary numbers*. Define $F_2(y)$ to be the function that has the same sequence of power series coefficients:

$$F(y) \;=\; 0 + 1\,y + 2\,y^2 + 4\,y^3 + 5\,y^4 + 8\,y^5 + 9\,y^6 + 10\,y^7 + 16\,y^8 + 17\,y^9 + 18\,y^{10} + \ldots \tag{36.11-7}$$

A slightly more general function $F_b(y)$ (which for $b = 2$ gives the power series above) can be computed by the iteration

$$\begin{aligned} L_0 &= 0, \quad R_0 = y, \quad l_0 = y, \quad r_0 = y && \text{(36.11-8a)} \\ A_0 &= 1, \quad B_0 = 1, \quad b = 2 && \text{(36.11-8b)} \\ A_{n+1} &= b\,B_n && \text{(36.11-8c)} \\ B_{n+1} &= b\,[B_n + r_n\,A_n] && \text{(36.11-8d)} \\ l_{n+1} &= r_n = y^{F_{n+1}} && \text{(36.11-8e)} \\ r_{n+1} &= r_n\,l_n = y^{F_{n+2}} && \text{(36.11-8f)} \\ L_{n+1} &= R_n && \text{(36.11-8g)} \\ R_{n+1} &= R_n + r_{n+1}\,[L_n + A_{n+1}] \quad \to F_b(y) && \text{(36.11-8h)} \end{aligned}$$

A pari/gp implementation is

```
1   ffb(y, b=2, N=13)=
2   { /* correct up to order fib(N+3)-1 */
3       local(t, yl, yr, L, R, Lp, Rp, Ri, Li);
4       L=0;  R=0+1*y;
5       Li=1;  Ri=1;
6       yl=y;  yr=y;
7       for (k=1, N,
8           Li*=b;  Ri*=b;
9           Lp=Ri;  Rp=Ri+yr*Li;  Li=Lp;  Ri=Rp;
10          t=yr;  yr*=yl;  yl=t;
11          Lp=R;  Rp=R+yr*(L+Li);  L=Lp;  R=Rp;
12      );
13      return( R )
14  }
```

The sequence of coefficients

```
1, 6, 14, 35, 90, 234, 611, 1598, 4182, 10947, 28658, 75026, 196419, 514230, ...
```

coincides (disregarding the initial one) with entries A032908 and A093467 of [240]. Let $B(x)$ be the function with power series coefficients equal to one if the exponent is a Fibbinary number and zero else:

$$B(x) := 1 + x + x^2 + x^4 + x^5 + x^8 + x^9 + x^{10} + x^{16} + x^{17} + x^{18} + x^{20} + \ldots \qquad \text{(36.11-9)}$$

Then a functional equation for $B(x)$ is (see entry A003714 of [240])

$$B(x) = x\,B(x^4) + B(x^2) \qquad \text{(36.11-10)}$$

We turn the relation into an recursion for the computation of $B(x)$ correct up to the term x^N:

```
1   fibbi(z, R)=
2   {
3       if ( z+R==0, return( 1+R ) );
4       return( z*fibbi(z^4,R) + fibbi(z^2,R) );
5   }
```

We check the functional relation:

```
? N=30; R=O(x^(N+1));  \\ R is used to truncate terms of order >N
? t=fibbi(x,R)
  1 + x + x^2 + x^4 + x^5 + x^8 + x^9 + x^10 + x^16 + x^17 + x^18 + x^20 + x^21 + O(x^31)
? t2=fibbi(x^2,R)
  1 + x^2 + x^4 + x^8 + x^10 + x^16 + x^18 + x^20 + O(x^31)
? t4=x*fibbi(x^4,R)
  x + x^5 + x^9 + x^17 + x^21 + O(x^32)
? t-(t4+t2)
  O(x^31)
```

36.11.2 Digit extract algorithms for the rabbit constant

The *spectrum* of a real number x is the sequence of integers $\lfloor k \cdot x \rfloor$ where $k = 1, 2, 3, \ldots$. As mentioned in [268], the spectrum of the golden ratio $g = \frac{\sqrt{5}+1}{2} = 1.61803\ldots$ gives the exponents of y where the series for $y\,A(y)$ has coefficient one:

```
1   bt(x, n=25)=
2   {
3       local(v);
4       v = vector(n);
5       for (k=1, n, v[k]=floor(x*k));
6       return ( v );
7   }
```

```
g=(sqrt(5)+1)/2
 1.6180339887498948482045868343656381177203091798057628 62
```

```
n=40;
bt(g, n)
  [1, 3, 4, 6, 8, 9, 11, 12, 14, 16, 17, 19, 21, 22, 24, 25, 27, 29,
  30, 32, 33, 35, 37, 38, 40, 42, 43, 45, 46, 48, 50, 51, 53, 55, 56,
  58, 59, 61, 63, 64]
```

```
t=taylor(y*fa(y),y)
  y + y^3 + y^4 + y^6 + y^8 + y^9 + y^11 + y^12 + y^14 + y^16 + y^17 +
  y^19 + y^21 + y^22 + y^24 + y^25 + y^27 + y^29 + y^30 + y^32 + y^33 +
  y^35 + y^37 + y^38 + y^40 + y^42 + y^43 + y^45 + y^46 + y^48 + y^50 +
  y^51 + y^53 + y^55 + y^56 + y^58 + y^59 + y^61 + y^63 + y^64 + O(y^66)
```

The sequence [1, 3, 4, 6, ...] of exponents where the coefficient equals one is sequence A000201 of [240]. There is a digit extract algorithm for the binary expansion of the rabbit constant. We use a binary search algorithm:

```
1   bts(x, k)=
2   { /* return 0 if k is not in the spectrum of x, else return index >=1  */
3       local(nlo, nhi, t);
4       if ( 0==k, return(0) );
5       t = 1 + ceil(k/x);  \\ floor(t*x)>=k
6       nlo = 1;  nhi = t;
7       while ( nlo!=nhi,
8           t = floor( (nlo+nhi)/2 );
9           if ( floor(t*x) < k,  nlo=t+1,  nhi=t);
10      );
11      if ( floor(nhi*x) == k,  return(nhi),  return (0));
12  }
```

```
g=(sqrt(5)+1)/2
for(k=1,65,if(bts(g,k),print1("1"),print1("0")));print();
  10110101101101011010110110101101101011010110110101101011011010110
```

The algorithm is very fast, we compute 1000 bits starting from position 1,000,000,000,000:

```
g=(sqrt(5)+1)/2
dd=10^12;  /* digits starting at position dd... */
for(k=dd,dd+1000,if(bts(g,k),print1("1"),print1("0")));print();
  11011010110110101101011011010110101101101011010110101101101011010110
  [--snip--]
  ***   last result computed in 236 ms.
```

The connection between the sequence of lowest Fibonacci bits and the rabbit constants allows even more. Subtracting the Fibonacci numbers > 1 until zero or one is reached, gives the complement of the rabbit sequence:

```
1   fpn=999;
2   vpv=vector(fpn, j, fibonacci(j+2)); /* vpv=[2,3,5,8,...] */
3   t=vpv[length(vpv)];  /* log(t)/log(10)== 208.8471.  OK for range up to >10^200 (!) */
4
5   flb(x)=
6   { /* return the lowest bit of the Fibonacci representation */
7       local(k, t);
8       k=bsearchgeq(x, vpv);
9       while ( k>0,
10          t = vpv[k];
11          if (x>=t,  x-=t);
12          k-- );
13      return ( x );
14  }
```

```
dd=0;
for(k=dd,dd+40,t=flb(k);print1(1-t))
     10110101101101011010110110101101101011010
/* 0.10110101101101011010110110101101101011010  rabbit constant */
```

The routine `bsearchgeq()` does a binary search (see section 3.2 on page 125) for the first element that is greater or equal to the element sought:

```
1   bsearchgeq(x, v)=
2   { /* return index of first element in v[] that is >=x, return 0 if x>max(v[])  */
3       local(nlo, nhi, t);
4       nlo = 1;  nhi = length(v);
5       while ( nlo!=nhi,
6           t = floor( (nlo+nhi)/2 );
7           if ( v[t] < x,  nlo=t+1,  nhi=t);
8       );
9       if ( v[nhi] >= x,  return(nhi),  return (0));
10  }
```

We compute the first 1000 bits starting from position 10^{100}:

```
dd=10^100-1;
for (k=dd, dd+1000, t=flb(k); print1(1-t))
  1101011011010110101101101101101101101101011011010
  [--snip--]
  ***   last result computed in 1,305 ms.
```

36.12 Iterations related to the Pell numbers

Replacement rules: simultaneously replace all zeros by one ($0 \to 1$) and all ones by one-one-zero ($1 \to 110$).

```
B0 =  0
B1 =  1
B2 =  110
B3 =  1101101
B4 =  11011011101101110
B5 =  1101101110110111011011101101110110111011011101101101
B --> 1101101110110111011011101101110110111011011101101110110111011011101101110...
```

Now the construction is $B_n = B_{n-1}.B_{n-1}.B_{n-2}$. The length of the n-th string is

$$p_n \quad = \quad 1,\ 1,\ 3,\ 7,\ 17,\ 41,\ 99,\ 239,\ \ldots \qquad p_k = 2p_{k-1} + p_{k-2}$$

This sequence is entry A001333 of [240], the numerators of the continued fraction of $\sqrt{2}$. The sequence `B` of zeros and ones is entry A080764. The Pell numbers are the first differences (and the denominators of the continued fraction of $\sqrt{2}$), sequence A000129:

$$0,\ 1,\ 2,\ 5,\ 12,\ 29,\ 70,\ 169,\ 408,\ 985,\ 2378,\ 5741,\ \ldots$$

Now define the function $B(y)$ by the iteration

$$\begin{aligned}
L_0 &= 1, \qquad R_0 = 1 + y, \qquad l_0 = y, \qquad r_0 = y && \text{(36.12-1a)}\\
l_{n+1} &= r_n && \text{(36.12-1b)}\\
r_{n+1} &= r_n^2\, l_n && \text{(36.12-1c)}\\
L_{n+1} &= R_n && \text{(36.12-1d)}\\
R_{n+1} &= R_n + r_{n+1}\, R_n + r_{n+1}^2\, L_n \quad \to B(y) && \text{(36.12-1e)}
\end{aligned}$$

After the n-th step the series in y is correct up to order p_n. That is, the order of convergence is $\sqrt{2}+1 \approx 2.4142$. Implementation in pari/gp:

```
1   fb(y, N=8)=
2   {
3       local(t, yr, yl, L, R, Lp, Rp);
4       L=1;  R=1+y;  yl=y;  yr=y;
5       for(k=1,N,
6         t=yr; yr*=yr*yl; yl=t;
7         Lp=R; Rp=R+yr*R+yr^2*L; L=Lp; R=Rp;
8       );
9       return( R )
10  }
```

We obtain the series

$$B(y) \;=\; 1 + y + y^3 + y^4 + y^6 + y^7 + y^8 + y^{10} + y^{11} + y^{13} + y^{14} + y^{15} + y^{17} + y^{18} + y^{20} + \ldots \quad (36.12\text{-}2)$$

Define the *Pell constant* B as $B = \frac{1}{2}\,B(\frac{1}{2})$, then

$$\begin{aligned} B &= 0.858267656461002055792260308433375148664905190083506778 6\ldots \quad (36.12\text{-}3)\\ [\text{base } 2] &= 0.110110111011011101101110110111011011101101110110111011011\ldots\\ [\text{CF}] &= [0, 1, 6, 18, 1032, 16777344, 288230376151842816,\\ &\quad 139379657490816394634598239204272161737932 8, \ldots] \end{aligned}$$

For the terms of the continued fraction we note

```
                                                6 ==   2^2 +  2^1
                                               18 ==   2^4 +  2^1
                                             1032 ==  2^10 +  2^3
                                         16777344 ==  2^24 +  2^7
                                288230376151842816 ==  2^58 + 2^17
139379657490816394634598239204272161737932 8 == 2^140 + 2^41
```

36.12.1 Pell palindromes

Define the function $P(y)$ by

$$\begin{aligned} L_0 &= 1, \qquad R_0 \;=\; 1 + y^2, \qquad l_0 \;=\; y, \qquad r_0 \;=\; y && (36.12\text{-}4a)\\ l_{n+1} &= r_n && (36.12\text{-}4b)\\ r_{n+1} &= r_n^2\, l_n && (36.12\text{-}4c)\\ L_{n+1} &= R_n && (36.12\text{-}4d)\\ R_{n+1} &= R_n + r_{n+1}\, L_n + r_{n+1}\, l_{n+1}\, R_n \quad \to P(y) && (36.12\text{-}4e) \end{aligned}$$

Note that R_0 is a palindrome and in relation 36.12-4e the combination of the parts gives a palindrome. For $R_0 = 1 + y + y^2$ the iteration computes $\frac{1}{1-y}$.

Define the *Pell palindromic constant* as $P = P(1/2)/2$, then

$$\begin{aligned} P &= 0.732160433063532837164590187177304465727298658960411239 0\ldots \quad (36.12\text{-}5)\\ [\text{base } 2] &= 0.101110110110111011011101101110110111011011101101110110111011011\ldots\\ [\text{CF}] &= [0, 1, 2, 1, 2, 1, 3, 17, 1, 7, 2063, 1, 63, 268437503, 1, 8191, 590295810358974087167,\\ &\quad 1, 1073741823, 3741444191567111470601433171759587488422774366535 67, 1, \ldots] \end{aligned}$$

By construction, the binary expansion is a palindrome up to lengths $1, 3, 7, 17, 41, 99, 239, \ldots$.

The sequence of zeros and ones in the binary expansion of P is entry A104521 of [240]. It can be obtained by the replacement rules $(0 \to 1)$ and $(1 \to 101)$:

```
P0 =  0
P1 =  1
P2 =  101
P3 =  1011101
P4 =  10111011011011101
P5 =  1011101101101110110111011011101101110110111011011101
P --> 1011101101101110110111011011101101110110111011011101101110110111011011101...
```

36.12.2 Pell representation

To obtain the Pell representation of a given number n, set $t = n$ repeatedly subtract from t the largest number $p_k \in \{1, 3, 7, 17, 41, 99, 239, \ldots\}$ that is not greater than t. Stop when $t = 0$. The number of times that p_k has been subtracted gives the k-th digit of the representation. The resulting digits are 0, 1, or 2. If the k-th digit equals 2 then the $(k-1)$-th digit will be zero.

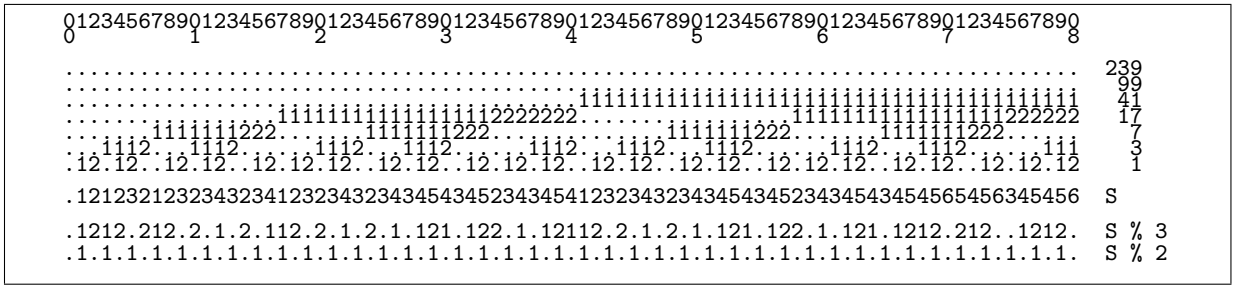

Figure 36.12-A: Pell representations of the numbers $0 \ldots 80$ (dots for zeros). The three lower lines are the sum of digits and the sum of digits modulo three and two.

The function $S(y)$ that has Taylor series coefficients corresponding to the sum of Pell digits

$$S(y) \;=\; 0 + 1y + 2y^2 + 1y^3 + 2y^4 + 3y^5 + 2y^6 + 1y^7 + 2y^8 + 3y^9 + 2y^{10} + 3y^{11} + 4y^{12} + \ldots \qquad (36.12\text{-}6)$$

can be computed via the iteration (see section 36.7 on page 726)

$$\begin{aligned}
L_0 &= 0, \qquad R_0 = 0 + y + 2\,y^2, \qquad l_0 = y, \qquad r_0 = y && (36.12\text{-}7a)\\
A_0 &= 1, \qquad B_0 = 1 + y + y^2 && (36.12\text{-}7b)\\
l_{n+1} &= r_n && (36.12\text{-}7c)\\
r_{n+1} &= r_n^2\, l_n && (36.12\text{-}7d)\\
L_{n+1} &= R_n && (36.12\text{-}7e)\\
R_{n+1} &= R_n + r_{n+1}\,(R_n + B_n) + r_{n+1}^2\,(L_n + 2\,A_n) \quad \to S(y) && (36.12\text{-}7f)\\
A_{n+1} &= B_n && (36.12\text{-}7g)\\
B_{n+1} &= B_n + r_{n+1}\,B_n + r_{n+1}^2\,A_n && (36.12\text{-}7h)
\end{aligned}$$

Implementation in pari/gp:

```
1   fs(y, N=8)=
2   {
3       local(t, yr, yl, L, R, Lp, Rp, Li, Ri);
4       L =  0;  R = 0+y+2*y^2;
5       Li = 1;  Ri = 1+y+y^2;
6       yl = y;  yr = y;
7       for(k=1,N,
8         t=yr; yr*=yr*yl; yl=t;
9         Lp=R;  Rp=R+yr*(R+Ri)+yr^2*(L+2*Li);  L=Lp; R=Rp;
10        Lp=Ri; Rp=Ri+yr*Ri+yr^2*Li;   Li=Lp; Ri=Rp;
11      );
12      return( R )
13  }
```

The series coefficients grow slowly, so the first few of them can nicely be displayed as

$$S\left(\frac{1}{10}\right) \;=\; 0.12123212323432341232343234345434523434541232343234345 \ldots \qquad (36.12\text{-}8)$$

36.12.3 Pell Gray code

Figure 36.12-B gives a Gray code for the Pell representations. The Gray code can be constructed recursively as shown in figure 36.12-C. In the algorithm each block is split into a left and a right part (indicated by the '^'. The next block is created by appending to the current block its reverse with ones on top and appending the left part with twos on top. The iteration can actually be started with a block of a single

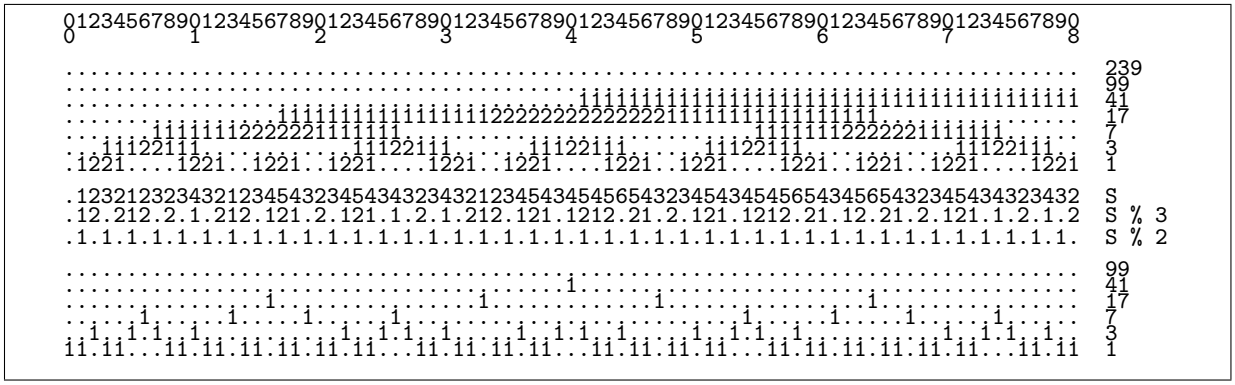

Figure 36.12-B: A Gray code for Pell representations. A dot is used for zero. The three following lines are the sum of digits and the sum of digits modulo three and two. The sequence is $0, 1, 2, 5, 4, 3, 6, 13, 10, 11, 12, 9, 8, 7, 14 \ldots$, the difference between successive elements is a Pell number. The lowest block gives the Pell representations of the (absolute) differences, the *Pell ruler function*.

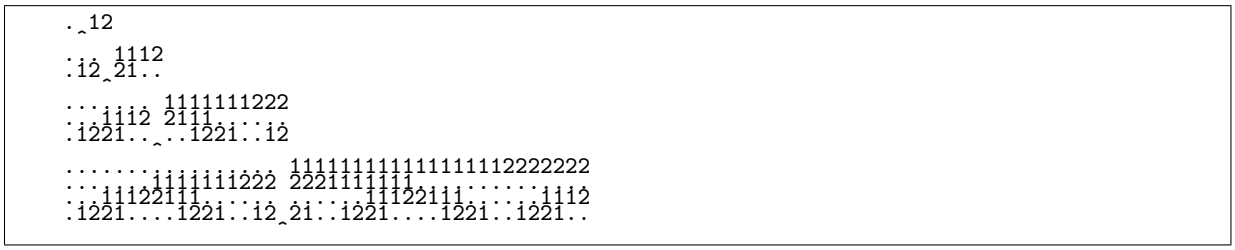

Figure 36.12-C: Construction for a Gray code for Pell representations.

zero (the left part being also a single zero). This is done in the following algorithm.

$$F_0 = 0, \quad F'_0 = 0, \quad B_0 = 0, \quad B'_0 = 0 \tag{36.12-9a}$$

$$I_0 = 1, \quad I'_0 = 1, \quad Y_0 = y, \quad Y'_0 = y \tag{36.12-9b}$$

$$b_n = 4^{n-1}, \quad c_n = 2\,b_n \tag{36.12-9c}$$

$$F_{n+1} = (F_n \qquad) + Y_n\,(B_n + b_n\,I_n) + \ Y_n^2\,(F'_n + c_n\,I'_n) \quad \to G_b(y) \tag{36.12-9d}$$

$$B_{n+1} = (B'_n + c_n\,I'_n) + Y'_n\,(F_n + b_n\,I_n) + Y'_n\,Y_n\,(B_n \qquad) \tag{36.12-9e}$$

$$I_{n+1} = I_n + Y_n\,I_n + Y_n^2\,I'_n \tag{36.12-9f}$$

$$Y_{n+1} = Y_n^2\,Y'_n \tag{36.12-9g}$$

$$F'_{n+1} = F_n, \quad B'_{n+1} = B_n \tag{36.12-9h}$$

$$I'_{n+1} = I_n, \quad Y'_{n+1} = Y_n \tag{36.12-9i}$$

Implementation in pari/gp:

```
1   pgr(y, N=11)= /* Pell Gray code */
2   {
3       local(iir, iil, yl, yr, Fl, F, Bl, B, b, c);
4       local(t, tf, tb);
5       /* correct up to order pell(N+1)-1 */
6       F=0;  Fl=0;   B=0;  Bl=0;
7       iil=1;  iir=1;  yl=y;  yr=y;
8       for(k=1, N,
9         b = 4^(k-1);  c = 2*b;  /* b = pell(k);*/
10        tf = (F         ) +  yr * (B + b*iir) + yr*yr * (Fl + c*iil);
11        tb = (Bl + c*iil) +  yl * (F + b*iir) + yl*yr * (B         );
12        Fl = F;   Bl = B;
13        F = tf;   B = tb;
14        t = iir;  iir += yr*(iir + yr*iil);  iil = t;
15        t = yr;  yr *= (yr*yl);  yl = t;
16      );
```

```
 ----- k = 0
yl = (y)  yr = (y)
iir = 1
iil = 1
F = 0
B = 0
 ----- k = 1    b=1
yl = (y)  yr = (y^3)
iir = (y^2 + y + 1)
iil = 1
F = (2 y^2 + y)
B = (y + 2)
 ----- k = 2    b=4
yl = (y^3)  yr = (y^7)
iir = (y^6 + y^5 + y^4 + y^3 + y^2 + y + 1)
iil = (y^2 + y + 1)
F = (8 y^6 + 4 y^5 + 5 y^4 + 6 y^3 + 2 y^2 + y)
B = (y^5 + 2 y^4 + 6 y^3 + 5 y^2 + 4 y + 8)
 ----- k = 3   b=16
yl = (y^7)  yr = (y^17)
iir = (y^16 + y^15 + y^14 + y^13 + y^12 + y^11 + y^10 + y^9
      + y^8 + y^7 + y^6 + y^5 + y^4 + y^3 + y^2 + y + 1)
iil = (y^6 + y^5 + y^4 + y^3 + y^2 + y + 1)
F = (34 y^16 + 33 y^15 + 32 y^14
      + 16 y^13 + 17 y^12 + 18 y^11 + 22 y^10 + 21 y^9 + 20 y^8 + 24 y^7
      + 8 y^6 + 4 y^5 + 5 y^4 + 6 y^3 + 2 y^2 + y)
B = (y^15 + 2 y^14 + 6 y^13 + 5 y^12 + 4 y^11 + 8 y^10
      + 24 y^9 + 20 y^8 + 21 y^7 + 22 y^6 + 18 y^5 + 17 y^4 + 16 y^3
      + 32 y^2 + 33 y + 34)
```

Figure 36.12-D: Quantities with the computation of the series related to the Pell Gray code.

```
17      return( F )
18  }
```

It is instructive to look at the variables in the first few steps of the iteration, see figure 36.12-D.

The Taylor series for $G_2(y)$ is

$$\begin{aligned} G_2(y) &= 0 + 1y + 2y^2 + 6y^3 + 5y^4 + 4y^5 + 8y^6 + 24y^7 + 20y^8 + 21y^9 + 22y^{10} + 18y^{11} + \\ &\quad +17y^{12} + 16y^{13} + 32y^{14} + 33y^{15} + 34y^{16} + 98y^{17} + 97y^{18} + 96y^{19} + 80y^{20} + 81y^{21} + \\ &\quad +82y^{22} + 86y^{23} + 85y^{24} + 84y^{25} + 88y^{26} + 72y^{27} + 68y^{28} + 69y^{29} + 70y^{30} + 66y^{31} + \dots \end{aligned} \tag{36.12-10}$$

The coefficients corresponds to the Pell representations interpreted as binary numbers, each Pell-digit occupying two bits (figure 36.12-B).

When relation 36.12-9c is changed to $b_n = P_n$ (indicated in the code, the function can be defined as `pell(k)=if(k<=1, 1, return(2*pell(k-1)+pell(k-2)));`) then the function $G_P(y)$ which has the Pell Gray code sequence as coefficients is computed:

$$\begin{aligned} G_P(y) &= 0 + 1y + 2y^2 + 5y^3 + 4y^4 + 3y^5 + 6y^6 + 13y^7 + 10y^8 + 11y^9 + 12y^{10} + 9y^{11} + \\ &\quad +8y^{12} + 7y^{13} + 14y^{14} + 15y^{15} + 16y^{16} + 33y^{17} + 32y^{18} + 31y^{19} + 24y^{20} + 25y^{21} + \\ &\quad +26y^{22} + 29y^{23} + 28y^{24} + 27y^{25} + 30y^{26} + 23y^{27} + 20y^{28} + 21y^{29} + 22y^{30} + 19y^{31} + \dots \end{aligned} \tag{36.12-11}$$

Section 12.4 on page 302 gives a recursive algorithm to compute the words of the Pell Gray code.

Define the *Pell Gray code constant* as

$$G_P = G_P\left(\frac{1}{2}\right) \tag{36.12-12a}$$

$$= 2.245567348365072195720956572438998819867495229140192012\dots \tag{36.12-12b}$$

$$[\text{base } 2] = 10.001111101101110110000000011100100011000110000100001111\dots \tag{36.12-12c}$$

$$[\text{CF}] = [2, 4, 13, 1, 5, 1, 1, 1, 27, 1, 9, 1, 3, 8, 1, 2, 1, 1, 3, 14, 1, 8, 1, 1, 6, 3, 1, 1, 1, 2, 1, 7, 210, 1, 1, 3, 2, 1, 1, 10, 1, 1, 6, 1, 1, 2, 1, 2, 1, 4, 6, 12, 1, \dots]$$

Setting $b_n = 1,\ c_n = 2$ in the algorithm gives a function whose series coefficients are the sum of Pell Gray code digits:

$$G_{[1,2]}\left(\frac{1}{10}\right) \quad = \quad 0.123212323432123454323454343234321234543454565432345434 5\ldots \quad (36.12\text{-}13)$$

Using $b_n = 1,\ c_n = 0$ counts the ones in the Pell Gray code:

$$G_{[1,0]}\left(\frac{1}{10}\right) \quad = \quad 0.101210123212101012123234321232121010121012321212323432 3\ldots \quad (36.12\text{-}14)$$

while $b_n = 0,\ c_n = 1$ counts the twos:

$$G_{[0,1]}\left(\frac{1}{10}\right) \quad = \quad 0.011001100110011221100110011001100112211221122110011001100 110\ldots \quad (36.12\text{-}15)$$

The continued fraction of this constant has very large terms:

```
[0, 90, 1, 8, 1, 5501100164, 8, 5, 3, 2, 19, 2, 1, 2, 2, 1, 1, 5, 1,
54, 6, 1, 5, 22, 2, 6, 2, 2, 1, 22445, 1, 45, 2, 2, 5, 1, 5, 1, 8,
5446094555544609444716727 4, 1, 5, 2, 2, 7, 1, 1, 1, 2, 1, 27, 2, 2,
1, 17, 1, 1, 1, 1, 4, 2, 1, 4, 3, 3, 3, 1, 3, 1, 2, 1, 1, 29, 1, ...]
```

Finally, $(1-y)\,G_P(y)$ gives the (signed) Pell ruler function, the first 100 series coefficients are:

```
0 +1 +1 +3 -1 -1 +3 +7 -3 +1 +1 -3 -1 -1 +7 +1 +1 +17 -1 -1 -7 +1 +1
+3 -1 -1 +3 -7 -3 +1 +1 -3 -1 -1 +17 +1 +1 +3 -1 -1 +3 +41 -3 +1 +1
-3 -1 -1 -17 +1 +1 +3 -1 -1 +3 +7 -3 +1 +1 -3 -1 -1 +7 +1 +1 -17 -1
-1 -7 +1 +1 +3 -1 -1 +3 -7 -3 +1 +1 -3 -1 -1 +41 +1 +1 +3 -1 -1 +3
+7 -3 +1 +1 -3 -1 -1 +7 +1 +1 +99 -1
```

Part V

Algorithms for finite fields

Chapter 37

Modular arithmetic and some number theory

We introduce some basic concepts of number theory. We start by implementing the arithmetical operations modulo m. Basic concepts of number theory, like the order of an element, quadratic residues and primitive roots are developed. Selected algorithms such as the Rabin-Miller compositeness test and several primality tests are presented. Some recreational material is included, such as multigrades and solutions of certain Diophantine equations. Finally we give the Cayley-Dickson construction construction for hypercomplex numbers and compute their multiplication tables.

Modular arithmetic and the concepts of number theory are fundamental for many areas like cryptography, error correcting codes, and digital signal processing.

37.1 Implementation of the arithmetic operations

We implement the basic operations of modular arithmetic: addition, subtraction, multiplication, powering, inversion and division.

37.1.1 Addition and subtraction

Addition and subtraction modulo m can easily be implemented as [FXT: mod/modarith.h]:

```
1   inline umod_t sub_mod(umod_t a, umod_t b, umod_t m)
2   {
3       if ( a>=b )  return  a - b;
4       else         return  m - b + a;
5   }
6
7   inline umod_t add_mod(umod_t a, umod_t b, umod_t m)
8   {
9       if ( 0==b )  return  a;
10      // return sub_mod(a, m-b, m);
11      b = m - b;
12      if ( a>=b )  return  a - b;
13      else         return  m - b + a;
14  }
```

The type `umod_t` is an unsigned 64-bit integer. Care has been taken to avoid any overflow of intermediate results. A 'set'-, increment-, decrement- and negation function will further be useful:

```
1   inline umod_t set_mod(umod_t x, umod_t m)
2   { if ( x>=m )  x %= m;  return x; }
3
4   inline umod_t incr_mod(umod_t a, umod_t m)
```

```
5   { a++; if ( a==m )  a = 0;   return a; }
6
7   inline umod_t decr_mod(umod_t a, umod_t m)
8   { if ( a==0 )  a = m - 1;  else a--;  return a; }
9
10  inline umod_t neg_mod(umod_t b, umod_t m)
11  { if ( 0==b )  return 0;  else  return m - b; }
```

Two addition tables for the moduli 13 and 9 are shown in figure 37.1-A

	0	1	2	3	4	5	6	7	8	9	10	11	12
0	0	1	2	3	4	5	6	7	8	9	10	11	12
1	1	2	3	4	5	6	7	8	9	10	11	12	0
2	2	3	4	5	6	7	8	9	10	11	12	0	1
3	3	4	5	6	7	8	9	10	11	12	0	1	2
4	4	5	6	7	8	9	10	11	12	0	1	2	3
5	5	6	7	8	9	10	11	12	0	1	2	3	4
6	6	7	8	9	10	11	12	0	1	2	3	4	5
7	7	8	9	10	11	12	0	1	2	3	4	5	6
8	8	9	10	11	12	0	1	2	3	4	5	6	7
9	9	10	11	12	0	1	2	3	4	5	6	7	8
10	10	11	12	0	1	2	3	4	5	6	7	8	9
11	11	12	0	1	2	3	4	5	6	7	8	9	10
12	12	0	1	2	3	4	5	6	7	8	9	10	11

	0	1	2	3	4	5	6	7	8
0	0	1	2	3	4	5	6	7	8
1	1	2	3	4	5	6	7	8	0
2	2	3	4	5	6	7	8	0	1
3	3	4	5	6	7	8	0	1	2
4	4	5	6	7	8	0	1	2	3
5	5	6	7	8	0	1	2	3	4
6	6	7	8	0	1	2	3	4	5
7	7	8	0	1	2	3	4	5	6
8	8	0	1	2	3	4	5	6	7

Figure 37.1-A: Addition modulo 13 (left) and modulo 9 (right).

37.1.2 Multiplication

Multiplication is a bit harder: if one would use something like

```
1   inline umod_t mul_mod(umod_t a, umod_t b, umod_t m)
2   {
3       return  (a * b) % m;
4   }
```

Then the modulus would be restricted to half of the word size.

One can use almost all bits for the modulus if the following trick is used. The technique also works if the product $a \cdot b$ does not fit into a machine integer.

Let $\langle x \rangle_y$ denote x modulo y, let $\lfloor x \rfloor$ denote the integer part of x. For $0 \leq a, b < m$:

$$a \cdot b \quad = \quad \left\lfloor \frac{a \cdot b}{m} \right\rfloor \cdot m + \langle a \cdot b \rangle_m \tag{37.1-1}$$

Rearranging and taking both sides modulo $z > m$ (where $z = 2^k$ on a k-bit machine):

$$\left\langle a \cdot b - \left\lfloor \frac{a \cdot b}{m} \right\rfloor \cdot m \right\rangle_z \quad = \quad \langle \langle a \cdot b \rangle_m \rangle_z \tag{37.1-2}$$

The right hand side equals $\langle a \cdot b \rangle_m$ because $m < z$.

$$\langle a \cdot b \rangle_m \quad = \quad \left\langle \langle a \cdot b \rangle_z - \left\langle \left\lfloor \frac{a \cdot b}{m} \right\rfloor \cdot m \right\rangle_z \right\rangle_z \tag{37.1-3}$$

The expression on the right can be translated into a few lines of C-code. The code given here assumes that one has 64-bit integer types `int64` (signed) and `uint64` (unsigned) and a floating point type with 64-bit mantissa, `float64` (typically `long double`).

```
uint64 mul_mod(uint64 a, uint64 b, uint64 m)
{
    uint64 y = (uint64)((float64)a*(float64)b/m+(float64)1/2);  // floor(a*b/m)
    y = y * m;           // m*floor(a*b/m) mod z
    uint64 x = a * b;    // a*b mod z
    uint64 r = x - y;    // a*b mod z - m*floor(a*b/m) mod z
    if ( (int64)r < 0 )  // normalization needed ?
    {
        r = r + m;
        y = y - 1;       // (a*b)/m  quotient, omit line if not needed
    }
    return  r;           // (a*b)%m  remnant
}
```

The technique uses the fact that integer multiplication computes the least significant bits of the result $\langle a \cdot b \rangle_z$ whereas float multiplication computes the most significant bits of the result. The above routine works if $0 <= a, b < m < 2^{63} = \frac{z}{2}$. The normalization is not necessary if $m < 2^{62} = \frac{z}{4}$.

When working with a fixed modulus the division by `p` may be replaced by a multiplication with the inverse modulus, that only needs to be computed once:

precompute: `float64 i = (float64)1/m;`

and replace the line `uint64 y = (uint64)((float64)a*(float64)b/m+(float64)1/2);`

by `uint64 y = (uint64)((float64)a*(float64)b*i+(float64)1/2);`

so any division inside the routine is avoided. Beware that the routine cannot be used for $m >= 2^{62}$: it very rarely fails for moduli of more than 62 bits, due to the additional error when inverting and multiplying as compared to dividing alone. An implementation in is [FXT: mod/modarith.h]:

```
inline umod_t mul_mod(umod_t a, umod_t b, umod_t m)
{
    umod_t x = a * b;
    umod_t y = m * (umod_t)( (ldouble)a * (ldouble)b/m + (ldouble)1/2 );
    umod_t r = x - y;
    if ( (smod_t)r < 0 )  r += m;
    return  r;
}
```

	0	1	2	3	4	5	6	7	8	9	10	11	12
0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	1	2	3	4	5	6	7	8	9	10	11	12
2	0	2	4	6	8	10	12	1	3	5	7	9	11
3	0	3	6	9	12	2	5	8	11	1	4	7	10
4	0	4	8	12	3	7	11	2	6	10	1	5	9
5	0	5	10	2	7	12	4	9	1	6	11	3	8
6	0	6	12	5	11	4	10	3	9	2	8	1	7
7	0	7	1	8	2	9	3	10	4	11	5	12	6
8	0	8	3	11	6	1	9	4	12	7	2	10	5
9	0	9	5	1	10	6	2	11	7	3	12	8	4
10	0	10	7	4	1	11	8	5	2	12	9	6	3
11	0	11	9	7	5	3	1	12	10	8	6	4	2
12	0	12	11	10	9	8	7	6	5	4	3	2	1

	0	1	2	3	4	5	6	7	8
0	0	0	0	0	0	0	0	0	0
1	0	1	2	3	4	5	6	7	8
2	0	2	4	6	8	1	3	5	7
3	0	3	6	0	3	6	0	3	6
4	0	4	8	3	7	2	6	1	5
5	0	5	1	6	2	7	3	8	4
6	0	6	3	0	6	3	0	6	3
7	0	7	5	3	1	8	6	4	2
8	0	8	7	6	5	4	3	2	1

Figure 37.1-B: Multiplication modulo 13 (left) and modulo 9 (right).

Two multiplication tables for the moduli 13 and 9 are shown in figure 37.1-B. Note that for the modulus 9 some products $a \cdot b$ are zero though neither of a or b is zero. The tables were computed with the program [FXT: mod/modarithtables-demo.cc].

For alternative multiplication (and reduction) techniques see [200, ch.14]. One method of great practical importance is the Montgomery multiplication described in [202].

37.1.3 Exponentiation

The algorithm used for exponentiation (powering) is the binary exponentiation algorithm shown in section 27.6 on page 555:

```
1   inline umod_t pow_mod(umod_t a, umod_t e, umod_t m)
2   // Right-to-left scan
3   {
4       if ( 0==e )  { return 1; }
5       else
6       {
7           umod_t z = a;
8           umod_t y = 1;
9           while ( 1 )
10          {
11              if ( e&1 )  y = mul_mod(y, z, m);  // y *= z;
12              e >>= 1;
13              if ( 0==e )  break;
14              z = sqr_mod(z, m);  // z *= z;
15          }
16          return  y;
17      }
18  }
```

37.1.4 Inversion and division

Subtraction is the inverse of addition. In order to subtract an element b from another element a one can add the *additive inverse* $-b := m - b$ to a. Each element has an additive inverse.

Division is the inverse of multiplication. In order to divide an element a by another element b one can multiply a by the *multiplicative inverse*. But not all elements have a multiplicative inverse, only those elements b that are coprime to the modulus m (that is, $\gcd(b, m) = 1$). These elements are called *invertible* (modulo m) or *units*. For a prime modulus all elements except zero are invertible.

The computation of the GCD uses the *Euclidean algorithm* [FXT: mod/gcd.h]:

```
1   template <typename Type>
2   Type gcd(Type a, Type b)
3   // Return greatest common divisor of a and b.
4   {
5       if ( a < b )  swap2(a, b);
6       if ( b==0 )  return a;
7       Type r;
8       do
9       {
10          r = a % b;
11          a = b;
12          b = r;
13      }
14      while ( r!=0 );
15      return a;
16  }
```

A variant of the algorithm that avoids most of the (expensive) modular reductions is called the *binary GCD algorithm* [FXT: mod/binarygcd.h]:

```
1   template <typename Type>
2   Type binary_ugcd(Type a, Type b)
3   // Return greatest common divisor of a and b.
4   // Version for unsigned types.
5   {
6       if ( a < b )  swap2(a, b);
7       if ( b==0 )  return a;
8
9       Type r = a % b;
10      a = b;
11      b = r;
12      if ( b==0 )  return a;
13
14      ulong k = 0;
15      while ( !((a|b)&1) )  // both even
```

```
16      {
17          k++;
18          a >>= 1;
19          b >>= 1;
20      }
21
22      while ( !(a&1) )  a >>= 1;
23      while ( !(b&1) )  b >>= 1;
24
25      while ( 1 )
26      {
27          if ( a==b )  return  a << k;
28
29          if ( a < b )  swap2(a, b);
30          Type t = (a-b) >> 1;  // t>0
31
32          while ( !(t&1) )  t >>= 1;
33          a = t;
34      }
35  }
```

The complexity of this algorithm for N-bit numbers is $O(N^2)$. We note that an $O(N \log(N))$ algorithm is given in [246].

The *least common multiple* (LCM) of two numbers is

$$\operatorname{lcm}(a,\, b) \;=\; \frac{a \cdot b}{\gcd(a,\, b)} \;=\; \left(\frac{a}{\gcd(a,\, b)}\right) \cdot b \tag{37.1-4}$$

The latter form avoids overflow when using integer types of fixed size.

For modular inversion one can use the extended Euclidean algorithm (EGCD), which for two integers a and b finds $d = \gcd(a, b)$ and u, v so that $a\,u + b\,v = d$. Applying EGCD to b and m where $\gcd(b, m) = 1$ one obtains u and v so that $m\,u + b\,v = 1$. That is: $b\,v \equiv 1 \pmod{m}$, so v is the inverse of b modulo m and $a/b := a\,b^{-1} = a\,v$.

The following code implements the EGCD algorithm as given in [168]:

```
1   template <typename Type>
2   Type egcd(Type u, Type v, Type &tu1, Type &tu2)
3   // Return u3 and set u1,v1 so that
4   //   gcd(u,v) == u3 == u*u1 + v*u2
5   // Type must be a signed type.
6   {
7       Type u1 = 1,  u2 = 0;
8       Type v1 = 0,  v3 = v;
9       Type u3 = u,  v2 = 1;
10      while ( v3!=0 )
11      {
12          Type q = u3 / v3;
13
14          Type t1 = u1 - v1 * q;
15          u1 = v1;  v1 = t1;
16
17          Type t3 = u3 - v3 * q;
18          u3 = v3;  v3 = t3;
19
20          Type t2 = u2 - v2 * q;
21          u2 = v2;  v2 = t2;
22      }
23      tu1 = u1;  tu2 = u2;
24      return u3;
25  }
```

Another algorithm for the computation of the modular inversion uses exponentiation. It is given only after the concept of the *order* of an element has been introduced (section 37.4 on page 761).

```
2^64-2^32+1 == 2^(32*2)-2^(32*1)+1
2^96-2^32+1 == 2^(32*3)-2^(32*1)+1
2^224-2^96+1 == 2^(32*7)-2^(32*3)+1
2^320-2^288+1 == 2^(32*10)-2^(32*9)+1
2^512-2^32+1 == 2^(32*16)-2^(32*1)+1
2^512-2^288+1 == 2^(32*16)-2^(32*9)+1
2^544-2^32+1 == 2^(32*17)-2^(32*1)+1
2^544-2^96+1 == 2^(32*17)-2^(32*3)+1
2^576-2^512+1 == 2^(64*9)-2^(64*8)+1
2^672-2^192+1 == 2^(32*21)-2^(32*6)+1
2^832-2^448+1 == 2^(64*13)-2^(64*7)+1
2^992-2^832+1 == 2^(32*31)-2^(32*26)+1
2^1088-2^608+1 == 2^(32*34)-2^(32*19)+1
2^1184-2^768+1 == 2^(32*37)-2^(32*24)+1
2^1376-2^32+1 == 2^(32*43)-2^(32*1)+1
2^1664-2^256+1 == 2^(128*13)-2^(128*2)+1
2^1856-2^1056+1 == 2^(32*58)-2^(32*33)+1
2^1920-2^384+1 == 2^(128*15)-2^(128*3)+1
2^1984-2^544+1 == 2^(32*62)-2^(32*17)+1
```

```
2^80-2^48+1 == 2^(16*5)-2^(16*3)+1
2^176-2^48+1 == 2^(16*11)-2^(16*3)+1
2^176-2^80+1 == 2^(16*11)-2^(16*5)+1
2^368-2^336+1 == 2^(16*23)-2^(16*21)+1
2^384-2^80+1 == 2^(16*24)-2^(16*5)+1
2^400-2^160+1 == 2^(16*25)-2^(16*10)+1
2^528-2^336+1 == 2^(16*33)-2^(16*21)+1
2^544-2^304+1 == 2^(16*34)-2^(16*19)+1
2^560-2^112+1 == 2^(16*35)-2^(16*7)+1
2^576-2^240+1 == 2^(16*36)-2^(16*15)+1
2^672-2^560+1 == 2^(16*42)-2^(16*35)+1
2^688-2^96+1 == 2^(16*43)-2^(16*6)+1
2^784-2^48+1 == 2^(16*49)-2^(16*3)+1
2^832-2^432+1 == 2^(16*52)-2^(16*27)+1
2^880-2^368+1 == 2^(16*55)-2^(16*23)+1
2^912-2^32+1 == 2^(16*57)-2^(16*2)+1
2^944-2^784+1 == 2^(16*59)-2^(16*49)+1
2^1008-2^144+1 == 2^(16*63)-2^(16*9)+1
2^1024-2^880+1 == 2^(16*64)-2^(16*55)+1
```

Figure 37.2-A: The complete list of primes of the form $p = x^k - x^j + 1$ where $x = 2^G$, $G = 2^i$, $G \geq 32$ and p up to 2048 bits (left), and the equivalent list for $x = 2^{16}$ and p up to 1024 bits (right).

```
? M=x^20-x^15+x^10-x^5+1;
? n=poldegree(M);
? P=sum(i=0,2*n-1,eval(Str("p_"i))*x^i)
  p_39*x^39 + p_38*x^38 + [--etc--] + p_3*x^3 + p_2*x^2 + p_1*x + p_0
? R=P%M;
? for(i=0,n-1,print("  ",eval(Str("r_"i))," = ",polcoeff(R,i)))
  r_0 = p_0 + (-p_20 - p_25)
  r_1 = p_1 + (-p_21 - p_26)
  r_2 = p_2 + (-p_22 - p_27)
  r_3 = p_3 + (-p_23 - p_28)
  r_4 = p_4 + (-p_24 - p_29)
  r_5 = p_5 + (p_20 - p_30)
  r_6 = p_6 + (p_21 - p_31)
  r_7 = p_7 + (p_22 - p_32)
  r_8 = p_8 + (p_23 - p_33)
  r_9 = p_9 + (p_24 - p_34)
  r_10 = p_10 + (-p_20 - p_35)
  r_11 = p_11 + (-p_21 - p_36)
  r_12 = p_12 + (-p_22 - p_37)
  r_13 = p_13 + (-p_23 - p_38)
  r_14 = p_14 + (-p_24 - p_39)
  r_15 = p_15 + p_20
  r_16 = p_16 + p_21
  r_17 = p_17 + p_22
  r_18 = p_18 + p_23
  r_19 = p_19 + p_24
```

Figure 37.2-B: Computation of the reduction rule for the 640-bit prime $Y_{50}(2^{32})$.

37.2 Modular reduction with structured primes

The modular reduction with Mersenne primes $M = 2^k - 1$ is especially easy: let u and v be in the range $0 \le u, v < M = 2^k - 1$, then with the non-reduced product written as $u\,v = 2^k\,r + s$ (where $0 \le r, s < M = 2^k - 1$) the reduction is simply $u\,v \equiv r + s \pmod{M}$.

A modular reduction algorithm that uses only shifts, additions and subtractions can also be found for *structured primes* (called *generalized Mersenne primes* in [242]). Let the modulus M be of the form

$$M \;=\; \sum_{i=0}^{n} m_i\,x^i \tag{37.2-1}$$

where $x = 2^k$ and the leading coefficient is equal to one: $m_n = 1$. For simplicity we further assume that $m_i = \pm 1$ and $m_{n-1} = -1$ (so that the numbers fit into n bits). The reduction algorithm can be found using polynomial arithmetic. Write the non-reduced product P as

$$P \;=\; \sum_{i=0}^{2n-1} p_i\,x^i \tag{37.2-2}$$

where $0 \le p_i < x$. Then the reduced product R is

$$R \;=\; \sum_{i=0}^{n-1} r_i\,x^i \;:=\; P \pmod{M} \tag{37.2-3}$$

where $0 \le r_i < x$. We determine the reduction rules for moduli of the form $x^k - x^j + 1$ (for $k = 3$ and $j = 2$, the rules are the last three lines):

```
? k=3;j=2;
? M=x^k-x^j+1
  x^3 - x^2 + 1
? n=poldegree(M);
? P=sum(i=0,2*n-1,eval(Str("p_"i))*x^i) \\ unreduced product
  p_5*x^5 + p_4*x^4 + p_3*x^3 + p_2*x^2 + p_1*x + p_0
? R=P%M;  \\ reduced product
? for(i=0,n-1,print("  ",eval(Str("r_"i))," = ",polcoeff(R,i)))
  r_0 = p_0 + (-p_3 + (-p_4 - p_5))
  r_1 = p_1 + (-p_4 - p_5)
  r_2 = p_2 + (p_3 + p_4)
```

A list of primes of the form $p = x^k - x^j + 1$ where $x = 2^G$, G a power of two and $G \ge 16$ is shown in figure 37.2-A. The equivalent list with i a multiple of 8 is given in [FXT: data/structured-primes-2k2j1.txt]. The primes allow radix-2 number theoretic transforms up to a length of x^j.

Structured primes that are evaluations of cyclotomic polynomials are given in section 37.11.4.7 on page 791. The reduction rule for the 640-bit prime $M = Y_{50}(2^{32})$ is shown in figure 37.2-B. There is a choice for the 'granularity' of the rule: the modulus also equals $Y_{10}(2^{5\cdot 32})$, so we can obtain the reduction rule for groups of five 32-bit words

```
? M=x^4-x^3+x^2-x^1+1;
  [--snip--]
? for(i=0,n-1,print("  ",eval(Str("r_"i))," = ",polcoeff(R,i)))
  r_0 = p_0 + (-p_4 - p_5)
  r_1 = p_1 + (p_4 - p_6)
  r_2 = p_2 + (-p_4 - p_7)
  r_3 = p_3 + p_4
```

The rule in terms of single words seems to be more appropriate as it allows for easier code generation.

37.3 The sieve of Eratosthenes

Several number theoretic algorithms can take advantage of a precomputed list of primes. A simple and quite efficient algorithm, called the *sieve of Eratosthenes* computes all primes up to a given limit. It uses

a tag-array where all entries ≥ 2 are initially marked as potential primes. The algorithm proceeds by searching for the next marked entry and deleting all multiples of it.

An implementation that uses the `bitarray` class (see section 4.6 on page 160) is given in [FXT: mod/eratosthenes-demo.cc]:

```
1   void
2   eratosthenes(bitarray &ba)
3   {
4       ba.set_all();
5       ba.clear(0);
6       ba.clear(1);
7       ulong n = ba.n_;
8       ulong k = 0;
9       while ( (k=ba.next_set(k+1)) < n )
10      {
11          for (ulong j=2, i=j*k;  i<n;  ++j, i=j*k)  ba.clear(i);
12      }
13  }
```

The program prints the resulting list of primes (code slightly simplified):

```
int
main(int argc, char **argv)
{
    ulong n = 100;
    NXARG(n, "Upper limit for prime search");

    bitarray ba(n);
    eratosthenes(ba);

    ulong k = 0;
    ulong ct = 0;
    while ( (k=ba.next_set(k+1)) < n )
    {
        ++ct;
        cout << " " << k;
    }
    cout << endl;

    cout << "Found " << ct << " primes below " << n << "." << endl;

    return 0;
}
```

The output for the default ($n = 100$) is:

```
 2 3 5 7 11 13 17 19 23 29 31 37 41 43 47 53 59 61 67 71 73 79 83 89 97
Found 25 primes below 100.
```

A little thought leads to a faster variant: when deleting the multiples $k \cdot p$ of the prime p from the list we only need to care about the values of k that are greater than all primes found so far. Further, values $k \cdot p$ containing only prime factors smaller than p have already been deleted. That is, we only need to delete the values $\{p^2, p^2 + p, p^2 + 2p, p^2 + 3p, \ldots\}$. If we further extract the loop for the prime 2 then for the odd primes, we need to delete only the values $\{p^2, p^2 + 2p, p^2 + 4p, \ldots\}$.

The implementation is

```
1   void
2   eratosthenes_opt(bitarray &ba)
3   {
4       ba.set_all();
5       ba.clear(0);
6       ba.clear(1);
7       ulong n = ba.n_;
8       for (ulong k=4; k<n; k+=2)  ba.clear(k);
9       ulong r = isqrt(n);
10      ulong k = 0;
11      while ( (k=ba.next_set(k+1)) < n )
12      {
13          if ( k > r )  break;
14          for (ulong j=k*k; j<n; j+=k*2)  ba.clear(j);
15      }
16  }
```

The routine is included in the demo, the second argument chooses whether the optimized routine is used.

When computing the primes up to a limit N then about N/p values are deleted after finding the prime p. If we slightly overestimate the computational work W by

$$W \approx N \sum_{p<N,\, prime} \frac{1}{p} \tag{37.3-1}$$

then $W \approx N \log(\log(N))$ which is almost linear. Practically, much of the time used with larger values of N is lost waiting for memory access. Thereby, further improvements should rather address machines specific optimizations than additional algorithmic refinements.

The described algorithmic improvement can be deduced from the series acceleration of the Lambert series

$$\sum_{k>0} d(k)\, x^k = \sum_{k>0} \frac{x^k}{1-x^k} = \sum_{k>0}\sum_{j>0} x^{kj} = \sum_{k>0} \frac{1+x^k}{1-x^k}\, x^{k^2} \tag{37.3-2}$$

where $d(k)$ is the sum of the divisors of k. We note a relation from [169, p.644, ex.27]:

$$\sum_{k>0} \frac{x^k}{1-x^k} = \sum_{k>0} k\, x^k \left(1-x^{k+1}\right)\left(1-x^{k+2}\right)\left(1-x^{k+3}\right)\ldots \tag{37.3-3}$$

One can save half of the space by recording only the odd primes. A C++ implementation of the modified algorithm is [FXT: `make_oddprime_bitarray()` in mod/eratosthenes.cc]. The corresponding table is created upon startup of programs linking the FXT-library. The data can be used to verify the primality of small numbers [FXT: `is_small_prime()` in mod/primes.cc]. The function `next_small_prime()` in the same file uses the data to return the next prime greater or equal to its argument or zero if the argument is too big.

37.4 The order of an element

```
     0  1  2  3  4  5  6  7  8  9 10 11 12 <--- exponent
-----------------------------------------   [order]
 0   1  0  0  0  0  0  0  0  0  0  0  0  0   [ --]
 1   1  1  1  1  1  1  1  1  1  1  1  1  1   [  1]
 2   1  2  4  8  3  6 12 11  9  5 10  7  1   [ 12]
 3   1  3  9  1  3  9  1  3  9  1  3  9  1   [  3]
 4   1  4  3 12  9 10  1  4  3 12  9 10  1   [  6]
 5   1  5 12  8  1  5 12  8  1  5 12  8  1   [  4]
 6   1  6 10  8  9  2 12  7  3  5  4 11  1   [ 12]
 7   1  7 10  5  9 11 12  6  3  8  4  2  1   [ 12]
 8   1  8 12  5  1  8 12  5  1  8 12  5  1   [  4]
 9   1  9  3  1  9  3  1  9  3  1  9  3  1   [  3]
10   1 10  9 12  3  4  1 10  9 12  3  4  1   [  6]
11   1 11  4  5  3  7 12  2  9  8 10  6  1   [ 12]
12   1 12  1 12  1 12  1 12  1 12  1 12  1   [  2]
```

Figure 37.4-A: Powers and orders modulo 13, the maximal order is $R(13) = 12 = \varphi(13)$.

The (multiplicative) *order* $r = \mathrm{ord}(a)$ of an element a is the smallest positive exponent so that $a^r = 1$. For elements that are not invertible ($\gcd(a, m) \neq 1$) the order is not defined. Figure 37.4-A shows the powers of all elements modulo the prime 13. The rightmost column gives the order of those elements that are invertible.

An element a whose r-th power equals one is called an *r-th root of unity*: $a^r = 1$. Modulo 9 both elements 2 and 4 are a 6th roots of unity, see figure 37.4-B.

	0	1	2	3	4	5	6	7	8	<--= exponent
										[order]
0	1	0	0	0	0	0	0	0	0	[--]
1	1	1	1	1	1	1	1	1	1	[1]
2	1	2	4	8	7	5	1	2	4	[6]
3	1	3	0	0	0	0	0	0	0	[--]
4	1	4	7	1	4	7	1	4	7	[3]
5	1	5	7	8	4	2	1	5	7	[6]
6	1	6	0	0	0	0	0	0	0	[--]
7	1	7	4	1	7	4	1	7	4	[3]
8	1	8	1	8	1	8	1	8	1	[2]

	0	1	2	3	4	5	6	<--= exponent
								[order]
1	1	1	1	1	1	1	1	[1]
2	1	2	4	8	7	5	1	[6]
4	1	4	7	1	4	7	1	[3]
5	1	5	7	8	4	2	1	[6]
7	1	7	4	1	7	4	1	[3]
8	1	8	1	8	1	8	1	[2]

Figure 37.4-B: Powers and orders modulo 9 (left), the maximal order is $R(9) = 6 = \varphi(9)$. The order modulo m is defined only for elements a where $\gcd(a, m) = 1$. The table of powers for the group of units $(\mathbb{Z}/9\mathbb{Z})^*$ (right) is obtained by dropping all elements for which the order is undefined.

If $a^r = 1$ but $a^x \neq 1$ for all $x < r$ then a is called a *primitive r-th root of unity*. Modulo 9 the element 2 is a primitive 6th root of unity; the element 4 is not, it is a primitive 3rd root of unity. An element of order r is an r-th primitive root of unity.

The *maximal order* $R(m)$ is the maximum of the orders of all elements for a fixed modulus m:

$$R(m) \quad := \quad \max_{a \in \mathbb{Z}/m\mathbb{Z}^*} (\mathrm{ord}(a)) \tag{37.4-1}$$

For prime modulus p the maximal order equals $R(p) = p - 1$. Where it cannot cause confusion we omit the argument to the maximal order in what follows.

An element of maximal order is a R-th primitive root of unity. Roots of unity of an order different from R are available only for the divisors d_i of R: if g is an element of maximal order R then g^{R/d_i} has order d_i (it is a primitive d_i-th root of unity):

$$\mathrm{ord}\left(g^{R/d_i}\right) \quad = \quad d_i \tag{37.4-2}$$

This is because $(g^{R/d_i})^{d_i} = g^R = 1$ and $(g^{R/d_i})^k \neq 1$ for $k < d_i$.

The factor by which the order of an element falls short of the maximal order is sometimes called the *index* of the element. Let i be the index and r the order, then $i \cdot r = R$.

The concept of the order comes from group theory. The invertible elements modulo m with multiplication form a group: the *multiplicative group*. The neutral element is one. The (multiplicative) order defined above is the order in this group, it tells us how often we have to multiply the element to one to obtain one. We restrict orders to positive values, else every element would have order zero.

With addition things are simpler, all elements with addition form a group with zero as neutral element: the *additive group*. The additive order of an element in this group tells us how often we have to add the element to zero to obtain zero. The additive order of the element a modulo m is simply $m/\gcd(a, m)$. All elements coprime to m (and especially 1 and -1) are generators of the additive group.

The maximal order R of all elements of a group is sometimes called the *exponent of the group*. Elements of maximal order are also called *primitive elements* of the group, *primitive roots* of the group, or *generators* of the group. In what follows we describe under which conditions the multiplicative group has generators.

37.5 Prime modulus: the field $\mathbb{Z}/p\mathbb{Z} = \mathbb{F}_p = \mathrm{GF}(p)$

If the modulus is a prime p then $\mathbb{Z}/p\mathbb{Z}$ is the field $\mathbb{F}_p = \mathrm{GF}(p)$: all elements except 0 have inverses and thereby division is possible in $\mathrm{GF}(p)$. The maximal order R equals $p - 1$. Elements of order R are called

primitive roots modulo p or *generators modulo p*.

If g is a generator, then every element in GF(p) different from 0 is equal to some power g^e ($1 \le e < p$) of g and its order is R/e. To test whether g is a primitive root we only need to check whether

$$g^{(p-1)/q_i} \quad \neq \quad 1 \bmod p \tag{37.5-1}$$

for all prime divisors of q_i of $p-1$. To find a primitive root use a simple search:

```
1   function primroot(p)
2   {
3       if p==2 then  return 1
4       f[] := distinct_prime_factors(p-1)
5       for r:=2 to p-1
6       {
7           x := TRUE
8           foreach q in f[]
9           {
10              if r**((p-1)/q)==1 then x:=FALSE
11          }
12          if x==TRUE then  return r
13      }
14      error("no primitive root found")  // p cannot be prime !
15  }
```

In practice the root is found after only a few tries. Note that the factorization of $p-1$ must be known. An element of order n in GF(p) is returned by the following function:

```
1   function element_of_order(n, p)
2   {
3       R := p-1  // maxorder
4       if (R/n)*n != R then  error("order n must divide maxorder p-1")
5       r := primroot(p)
6       x := r**(R/n)
7       return x
8   }
```

37.6 Composite modulus: the ring $\mathbb{Z}/m\mathbb{Z}$

n:	$\varphi(n)$	n:	$\varphi(n)$	n:	$\varphi(n)$	n:	$\varphi(n)$	n:	$\varphi(n)$	n:	$\varphi(n)$	n:	$\varphi(n)$	n:	$\varphi(n)$
1:	1	13:	12	25:	20	37:	36	49:	42	61:	60	73:	72	85:	64
2:	1	14:	6	26:	12	38:	18	50:	20	62:	30	74:	36	86:	42
3:	2	15:	8	27:	18	39:	24	51:	32	63:	36	75:	40	87:	56
4:	2	16:	8	28:	12	40:	16	52:	24	64:	32	76:	36	88:	40
5:	4	17:	16	29:	28	41:	40	53:	52	65:	48	77:	60	89:	88
6:	2	18:	6	30:	8	42:	12	54:	18	66:	20	78:	24	90:	24
7:	6	19:	18	31:	30	43:	42	55:	40	67:	66	79:	78	91:	72
8:	4	20:	8	32:	16	44:	20	56:	24	68:	32	80:	32	92:	44
9:	6	21:	12	33:	20	45:	24	57:	36	69:	44	81:	54	93:	60
10:	4	22:	10	34:	16	46:	22	58:	28	70:	24	82:	40	94:	46
11:	10	23:	22	35:	24	47:	46	59:	58	71:	70	83:	82	95:	72
12:	4	24:	8	36:	12	48:	16	60:	16	72:	24	84:	24	96:	32

Figure 37.6-A: Values of $\varphi(n)$, the number of integers less than n and coprime to n, for $n \le 96$.

In what follows we will need the function φ, the *totient function* (or *Euler's totient function*). The function $\varphi(m)$ counts the number of integers coprime to and less than m:

$$\varphi(m) \quad := \quad \sum_{\substack{1 \le k < m \\ \gcd(k,m)=1}} 1 \tag{37.6-1}$$

The sequence of values $\varphi(n)$ is entry A000010 in [240]. The values of $\varphi(n)$ for $n \leq 96$ are shown in figure 37.6-A. For $m = p$ prime one has $\varphi(p) = p - 1$. For m composite $\varphi(m)$ is always less than $m - 1$. For $m = p^k$ a prime power

$$\varphi(p^k) \quad = \quad p^k - p^{k-1} \; = \; p^{k-1}\,(p-1) \tag{37.6-2}$$

The totient function is a *multiplicative* function: one has $\varphi(x_1\,x_2) = \varphi(x_1)\,\varphi(x_2)$ for coprime x_1, x_2, that is, $\gcd(x_1, x_2) = 1$ (but x_1 and x_2 are not required to be primes). Thereby, if p_i are distinct primes in the factorization of n, then

$$\varphi(n) \quad = \quad \prod_i \varphi(p_i^{e_i}) \qquad \text{where} \quad n = \prod_i p_k^{e_i} \tag{37.6-3}$$

An alternative expression for $\varphi(n)$ is

$$\varphi(n) \quad = \quad n \prod_{p_i} \left(1 - \frac{1}{p_i}\right) \qquad \text{where} \quad n = \prod_i p_k^{e_i} \tag{37.6-4}$$

We note a generalization: the number of s-element sets of numbers $\leq n$ whose greatest common divisor is coprime to n equals

$$\varphi_s(n) \quad = \quad n^s \prod_{p_i} \left(1 - \frac{1}{p_i^s}\right) \qquad \text{where} \quad n = \prod_i p_i^{e_i} \tag{37.6-5}$$

Pseudo code to compute $\varphi(m)$ for arbitrary m:

```
1   function euler_phi(m)
2   {
3       {n, p[], x[]} := factorization(m)  // m==product(i=0..n-1, p[i]**x[i])
4       ph := 1
5       for i:=0 to n-1
6       {
7           k = := x[i]  // exponent
8           ph := ph * (p[i]**(k-1)) * (p[i]-1)  // ==ph * euler_phi(p[i]**x[i])
9       }
10  }
```

The multiplicative group (consisting of the invertible elements, or *units*) is denoted by $(\mathbb{Z}/m\mathbb{Z})^*$. The size of the group $(\mathbb{Z}/m\mathbb{Z})^*$ equals the number of units:

$$\left|(\mathbb{Z}/m\mathbb{Z})^*\right| \quad = \quad \varphi(m) \tag{37.6-6}$$

If m factorizes as $m = 2^{e_0} \cdot p_1^{e_1} \cdot \ldots \cdot p_q^{e_q}$ where p_i are pairwise distinct primes then

$$\left|(\mathbb{Z}/m\mathbb{Z})^*\right| \quad = \quad \varphi(2^{e_0}) \cdot \varphi(p_1^{e_1}) \cdot \ldots \cdot \varphi(p_q^{e_q}) \tag{37.6-7}$$

Further, the group $(\mathbb{Z}/m\mathbb{Z})^*$ is isomorphic to the direct product of the multiplicative groups modulo the prime powers:

$$(\mathbb{Z}/m\mathbb{Z})^* \quad \sim \quad (\mathbb{Z}/2^{e_0}\mathbb{Z}) \times (\mathbb{Z}/p_1^{e_1}\mathbb{Z}) \times \cdots \times \left(\mathbb{Z}/p_q^{e_q}\mathbb{Z}\right) \tag{37.6-8}$$

That is, instead of working modulo m we can do all computations modulo all prime powers in parallel. The Chinese remainder theorem (section 37.7 on page 769) tells us how to find the element modulo m given the results modulo the prime powers. The other direction is simply modular reduction.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	
0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	[--]
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	[1]
2	1	2	4	8	1	2	4	8	1	2	4	8	1	2	4	[4]
3	1	3	9	12	6	3	9	12	6	3	9	12	6	3	9	[--]
4	1	4	1	4	1	4	1	4	1	4	1	4	1	4	1	[2]
5	1	5	10	5	10	5	10	5	10	5	10	5	10	5	10	[--]
6	1	6	6	6	6	6	6	6	6	6	6	6	6	6	6	[--]
7	1	7	4	13	1	7	4	13	1	7	4	13	1	7	4	[4]
8	1	8	4	2	1	8	4	2	1	8	4	2	1	8	4	[4]
9	1	9	6	9	6	9	6	9	6	9	6	9	6	9	6	[--]
10	1	10	10	10	10	10	10	10	10	10	10	10	10	10	10	[--]
11	1	11	1	11	1	11	1	11	1	11	1	11	1	11	1	[2]
12	1	12	9	3	6	12	9	3	6	12	9	3	6	12	9	[--]
13	1	13	4	7	1	13	4	7	1	13	4	7	1	13	4	[4]
14	1	14	1	14	1	14	1	14	1	14	1	14	1	14	1	[2]

	0	1	2	3	4	
1	1	1	1	1	1	[1]
2	1	2	4	8	1	[4]
4	1	4	1	4	1	[2]
7	1	7	4	13	1	[4]
8	1	8	4	2	1	[4]
11	1	11	1	11	1	[2]
13	1	13	4	7	1	[4]
14	1	14	1	14	1	[2]

Figure 37.6-B: Powers and orders modulo 15 (left). The ring $\mathbb{Z}/15\mathbb{Z}$ is noncyclic: there are $\varphi(15) = 8$ invertible elements but no element generates all of them as the maximal order is $R(15) = 4 < \varphi(15)$. The table of powers for the group of units $(\mathbb{Z}/15\mathbb{Z})^*$ (right) is obtained by dropping all elements that are not invertible.

37.6.1 Cyclic and noncyclic rings

If the maximal order R if equal to $|(\mathbb{Z}/m\mathbb{Z})^*| = \varphi(m)$ then the ring $(\mathbb{Z}/m\mathbb{Z})^*$ is then called a *cyclic ring*, else we call the ring *noncyclic*. The term cyclic reflects that in those rings the powers of any element of maximal order 'cycle through' all elements of $(\mathbb{Z}/m\mathbb{Z})^*$. An element of maximal order in cyclic rings is also called a *generator* as its powers 'generate' all elements. Note that the terms cyclic and noncyclic refer to the muyltiplicative group. As the additive group is always cyclic (one is always a generator) there is no ambiguity when speaking of a cyclic ring.

Figure 37.6-B shows the powers and orders of the noncyclic ring $\mathbb{Z}/15\mathbb{Z}$ where no element generates all units. The rings $\mathbb{Z}/13\mathbb{Z}$ and $\mathbb{Z}/9\mathbb{Z}$ are cyclic, see figure 37.4-A on page 761, and figure 37.4-B.

For prime modulus $m = p$ the group $(\mathbb{Z}/m\mathbb{Z})^*$ contains all nonzero elements and any element of maximal order is a generator of the group.

For m a power p^k of an odd prime p the maximal order R in $(\mathbb{Z}/m\mathbb{Z})^*$ is

$$R(p^k) \quad = \quad \varphi(p^k) \tag{37.6-9}$$

For m a power of two an irregularity occurs:

$$R(2^k) \quad = \quad \begin{cases} 1 & \text{for } k = 1 \\ 2 & \text{for } k = 2 \\ 2^{k-2} & \text{for } k \geq 3 \end{cases} \tag{37.6-10}$$

That is, for powers of two greater than 4 the maximal order falls short from $\varphi(2^k) = 2^{k-1}$ by a factor of 2. For the general modulus $m = 2^{k_0} \cdot p_1^{k_1} \cdot \ldots \cdot p_q^{k_q}$ the maximal order is

$$R(m) \quad = \quad \text{lcm}\left(R(2^{k_0}),\, R(p_1^{k_1}),\, \ldots,\, R(p_q^{k_q})\right) \tag{37.6-11}$$

where lcm denotes the least common multiple. The maximal order $R(m)$ of an element in $\mathbb{Z}/m\mathbb{Z}$ can be computed as follows:

```
function maxorder(m)
{
    {n, p[], k[]} := factorization(m)  // m==product(i=0..n-1,p[i]**k[i])

    R := 1
    for i:=0 to n-1
    {
        t := euler_phi_pp(p[i], k[i])  // ==euler_phi(p[i]**k[i])
        if p[i]==2 AND k[i]>=3 then  t := t / 2
        R := lcm(R, t)
    }

    return R
}
```

Now we can see for which moduli m the ring $(\mathbb{Z}/m\mathbb{Z})^*$ will be cyclic (the maximal order equals the number of units $\varphi(m)$):

$$(\mathbb{Z}/m\mathbb{Z})^* \quad \text{is cyclic for} \quad m = 2,\ 4,\ p^k,\ 2 \cdot p^k \qquad \text{where } p \text{ is an odd prime} \tag{37.6-12}$$

If the factorization of m contains two different odd primes p_a and p_b then

$$R(m) \quad = \quad \text{lcm}(\dots, \varphi(p_a), \dots, \varphi(p_b), \dots)$$

is at least by a factor of two smaller than

$$\varphi(m) \quad = \quad \dots \cdot \varphi(p_a) \cdot \dots \cdot \varphi(p_b) \cdot \dots$$

because both $\varphi(p_a)$ and $\varphi(p_b)$ are even. Thereby $(\mathbb{Z}/m\mathbb{Z})^*$ cannot be cyclic in that case. The same argument holds for $m = 2^{k_0} \cdot p^k$ if $k_0 > 1$. For $m = 2^k$ the ring $(\mathbb{Z}/m\mathbb{Z})^*$ is cyclic only for $k = 1$ and $k = 2$ because of the mentioned irregularity of powers of two (relation 37.6-10).

Pseudo code for a function that returns the order of a given element x in $(\mathbb{Z}/m\mathbb{Z})^*$:

```
function order(x, m)
{
    if gcd(x,m)!=1 then  return 0  // x not a unit
    h := euler_phi(m)  // number of elements of ring of units
    e := h
    {n, p[], k[]} := factorization(h)  // h==product(i=0..n-1,p[i]**k[i])

    for i:=0 to n-1
    {
        f := p[i]**k[i]
        e := e / f
        g1 := x**e mod m
        while g1!=1
        {
            g1 := g1**p[i] mod m
            e := e * p[i]
        }
    }

    return e
}
```

Pseudo code for a function that returns an element x in $(\mathbb{Z}/m\mathbb{Z})^*$ of maximal order:

```
function maxorder_element(m)
{
    R := maxorder(m)
    for x:=1 to m-1
    {
        if order(x, m)==R then  return x
    }
    // never reached
}
```

Again, while the function does a simple search it is efficient in practice. For prime m the function returns a primitive root. A C++ implementation is [FXT: `maxorder_element_mod()` in mod/maxorder.cc]. Note that for noncyclic rings the returned element does not necessarily have maximal order modulo all factors of the modulus. We list all elements of $(\mathbb{Z}/15\mathbb{Z})^*$ together with their orders modulo 15, 3, and 5:

```
 1:  r=1  r3=1  r5=1
 2:  r=4  r3=2  r5=4
 4:  r=2  r3=1  r5=2
 7:  r=4  r3=1  r5=4  <--=
 8:  r=4  r3=2  r5=4
11:  r=2  r3=2  r5=1
13:  r=4  r3=1  r5=4  <--=
14:  r=2  r3=2  r5=2
```

The two elements marked with an arrow have maximal order modulo 15 but not modulo 3. An element of maximal order modulo all factors of a composite modulus (equivalently, maximal order in all subgroups) can be found by computing a generator for all cyclic subgroups and applying the Chinese remainder algorithm given in section 37.7 on page 769.

37.6.2 Generators in cyclic rings

Let G be the set of all generators in a cyclic ring modulo n. Then the number of generators is given by

$$|G| \quad = \quad \varphi(\varphi(n)) \tag{37.6-13}$$

Let g be a generator, the g^k is a generator exactly if $\gcd(k, \varphi(n)) = 1$, and there are $\varphi(\varphi(n))$ numbers k that are coprime to $\varphi(n)$.

Let g be a generator modulo a prime p. Then g is a generator modulo $2\,p^k$ for all $k \geq 1$ if g is odd. If g is even then $g + p^k$ is a generator modulo $2\,p^k$.

Further, g is a generator modulo p^k if $g^{p-1} \bmod p^2 \neq 1$. The only primes below $2^{36} \approx 68 \cdot 10^9$ for which the smallest primitive root is not the a generator modulo p^2 are 2, 40487 and 6692367337. Such primes are called *non-generous primes*, see entry A055578 of [240].

The only known primes p below $32 \cdot 10^{12}$ where $2^{p-1} = 1$ modulo p^2 are 1093 and 3511 (such primes are called *Wieferich primes*, see entry A001220 of [240]). As 2 is not a generator modulo either of the two we see that whenever 2 is a generator modulo $p < 32 \cdot 10^{12}$ then it is also a generator modulo p^k for all $k > 1$.

37.6.3 Generators in noncyclic rings

When the ring is cyclic an element of maximal order generates all invertible elements. With noncyclic rings one needs more than one generator. Pari/gp's function `znstar()` gives the complete information about the multiplicative group of units. The help text reads:

```
znstar(n): 3-component vector v, giving the structure of (Z/nZ)^*.
 v[1] is the order (i.e. eulerphi(n)),
 v[2] is a vector of cyclic components, and
 v[3] is a vector giving the corresponding  generators.
```

Its output for $2 \leq n \leq 25$ is shown in figure 37.6-C.

The ring is cyclic when there is just one generator. In general, when `znstar(n)` returns

$$[\varphi,\ [r_1, r_2, \ldots, r_k],\ [g_1, g_2, \ldots, g_k]] \tag{37.6-14}$$

then the φ invertible elements u are of the form

$$u \quad = \quad g_1^{e_1}\, g_2^{e_2}\, \ldots\, g_k^{e_k} \tag{37.6-15}$$

where $0 \leq e_i < r_i$ for $1 \leq i \leq k$. For example, with $n = 15$:

```
? for(n=2,25,print(n," ",znstar(n)))
     2 [1,  [],   []]    /* read:  [1, [1], [Mod(1,2)]] */
     3 [2,  [2],  [Mod(2, 3)]]
     4 [2,  [2],  [Mod(3, 4)]]
     5 [4,  [4],  [Mod(2, 5)]]
     6 [2,  [2],  [Mod(5, 6)]]
     7 [6,  [6],  [Mod(3, 7)]]
     8 [4,  [2, 2],  [Mod(5, 8), Mod(3, 8)]]
     9 [6,  [6],  [Mod(2, 9)]]
    10 [4,  [4],  [Mod(7, 10)]]
    11 [10, [10],  [Mod(2, 11)]]
    12 [4,  [2, 2],  [Mod(7, 12), Mod(5, 12)]]
    13 [12, [12],  [Mod(2, 13)]]
    14 [6,  [6],  [Mod(3, 14)]]
    15 [8,  [4, 2],  [Mod(8, 15), Mod(11, 15)]]
    16 [8,  [4, 2],  [Mod(5, 16), Mod(7, 16)]]
    17 [16, [16],  [Mod(3, 17)]]
    18 [6,  [6],  [Mod(11, 18)]]
    19 [18, [18],  [Mod(2, 19)]]
    20 [8,  [4, 2],  [Mod(3, 20), Mod(11, 20)]]
    21 [12, [6, 2],  [Mod(5, 21), Mod(8, 21)]]
    22 [10, [10],  [Mod(13, 22)]]
    23 [22, [22],  [Mod(5, 23)]]
    24 [8,  [2, 2, 2],  [Mod(13, 24), Mod(19, 24), Mod(17, 24)]]
    25 [20, [20],  [Mod(2, 25)]]
```

Figure 37.6-C: Structure of the multiplicative group modulo n for $2 \leq n \leq 25$, as reported by pari/gp's function `znstar()`.

```
? znstar(15)
  [8, [4, 2], [Mod(8, 15), Mod(11, 15)]]
? g1=Mod(8, 15); g2=Mod(11,15);
? for(e1=0,4-1,for(e2=0,2-1,print(e1," ",e2,"  ",g1^e1*g2^e2)))
  0 0  Mod(1, 15)
  0 1  Mod(11, 15)
  1 0  Mod(8, 15)
  1 1  Mod(13, 15)
  2 0  Mod(4, 15)
  2 1  Mod(14, 15)
  3 0  Mod(2, 15)
  3 1  Mod(7, 15)
```

Generators modulo $n = 2^k$

The ring modulo $n = 2^k$ is cyclic only for $k \leq 2$:

```
? for(i=1,6,print(i,": ",znstar(2^i)))
   1: [1, [], []]
   2: [2, [2], [Mod(3, 4)]]
   3: [4, [2, 2], [Mod(5, 8), Mod(3, 8)]]
   4: [8, [4, 2], [Mod(5, 16), Mod(7, 16)]]
   5: [16, [8, 2], [Mod(5, 32), Mod(15, 32)]]
   6: [32, [16, 2], [Mod(5, 64), Mod(31, 64)]]
```

For $k \geq 3$ the multiplicative group is generated by the two elements 5 and -1.

37.6.4 Inversion by exponentiation

For a unit u of order $r = \mathrm{ord}(u)$ one has $u^r = 1$. As r divides the maximal order R also $u^R = 1$ holds and thereby $u^{R-1} \cdot u = 1$. That is, the inverse of any invertible element u equals u to the $(R-1)$-st power:

$$u^{-1} \quad = \quad u^{R-1} \tag{37.6-16}$$

In fact, one has also $u^{-1} = u^{\varphi(m)-1}$ which may involve slightly more work if the ring is noncyclic.

37.7 The Chinese Remainder Theorem (CRT)

Let m_1, m_2, ..., m_f be pairwise coprime (that is, $\gcd(m_i, m_j) = 1$ for all $i \neq j$). If $x \equiv x_i \pmod{m_i}$ for $i = 1, 2, \ldots, f$ then x is unique modulo the product $M = m_1 \cdot m_2 \cdots m_f$. This is the *Chinese remainder theorem* (CRT). Note that it is not assumed that any of the m_i is prime.

The theorem tells us that a computation modulo a composite number M can be split into separate computations modulo the coprime factors of M. To evaluate a function $y := f(x) \bmod M$ where $M = m_1 \cdot m_2$ (with $\gcd(m_1, m_2) = 1$), proceed as follows

1. Splitting: compute $x_1 = x \bmod m_1$ and $x_2 = x \bmod m_2$.
2. Separate computations: compute $y_1 := f(x_1) \bmod m_1$ and $y_2 := f(x_2) \bmod m_2$
3. Recombination: compute y from y_1 and y_2 using the CRT

For example, when computing the exact convolution of a long sequence via number theoretic transforms (see section 25.3 on page 532) the largest term of the result must be smaller than the modulus. Assume that (efficient) modular arithmetic is available for moduli of at most word size. Now choose several coprime moduli whose product M is greater than the largest element of the result. Compute the transforms separately and only at the very end compute, via the CRT, the result modulo M.

37.7.1 Efficient computation

For two moduli m_1, m_2 compute x with $x \equiv x_1 \pmod{m_1}$ and $x \equiv x_2 \pmod{m_2}$ as suggested by the following pseudo code:

```
function crt2(x1, m1, x2, m2)
{
    c := m1**(-1) mod m2    // inverse of m1 modulo m2

    s := ((x2-x1)*c) mod m2

    return  x1 + s * m1
}
```

For repeated CRT calculations with the same moduli one will use precomputed values $c = m_1^{-1} \bmod m_2$. With more than two moduli use the above algorithm repeatedly. Pseudo code to perform the CRT for several moduli:

```
function crt(x[0,...,f-1], m[0,...,f-1], f)
{
    x1 := x[0]
    m1 := m[0]
    i := 1
    do
    {
        x2 := x[i]
        m2 := m[i]
        x1 := crt2(x1, m1, x2, m2)
        m1 := m1 * m2
        i := i + 1
    }
    while i < f
    return x1
}
```

A C++ implementation is given in [FXT: mod/chinese.cc]:

```
umod_t
chinese(const umod_t *x, const factorization &f)
// Return R modulo M where:
//   f[] is the factorization of M,
//   x[] := R modulo the prime powers of f[].
{
    const int n = f.nprimes();
    // (omitted test that gcd(m_0,...,m_{n-1})=1 )

    const umod_t M = f.product();
    umod_t R = 0;
    for (int i=0; i<n; ++i)
```

```
13      {
14          // Ti = prod(mk)  (where k!=i);  Ti==M/mi:
15          const umod_t Ti = M / f.primepow(i);  // exact division
16
17          // ci = 1 / Ti:
18          umod_t ci = inv_modpp(Ti, f.prime(i), f.exponent(i));
19          // here:  0 <= ci < mi
20
21          // Xi = x[i] * ci * Ti:
22          umod_t Xi = ci * Ti;  // 0 <= Xi < M
23          Xi = mul_mod(Xi, x[i], M);
24
25          // add Xi to result:
26          R = add_mod(R, Xi, M);
27      }
28
29      return R;
30  }
```

37.7.2 The underlying construction

We derive the algorithm for CRT recombination from a construction for k coprime moduli. Define T_i as

$$T_i \;:=\; \prod_{k\neq i} m_k \tag{37.7-1}$$

and c_i as

$$c_i \;:=\; T_i^{-1} \mod m_i \tag{37.7-2}$$

Then for X_i defined as

$$X_i \;:=\; x_i\, c_i\, T_i \tag{37.7-3}$$

one has

$$X_i \mod m_j \;=\; \begin{cases} x_i & \text{for} \quad j=i \\ 0 & \text{else} \end{cases} \tag{37.7-4}$$

Therefore

$$x \;:=\; \sum_k X_k \;=\; x_i \mod m_i \tag{37.7-5}$$

For the special case of two moduli m_1, m_2 one has

$$T_1 \;=\; m_2, \qquad T_2 \;=\; m_1 \tag{37.7-6a}$$

$$c_1 \;=\; m_2^{-1} \mod m_1, \qquad c_2 \;=\; m_1^{-1} \mod m_2 \tag{37.7-6b}$$

The quantities are related by

$$c_1\, m_2 + c_2\, m_1 \;=\; 1 \tag{37.7-7}$$

and

$$x \;=\; \sum_k X_k \;=\; x_1\, c_1\, T_1 + x_2\, c_2\, T_2 \tag{37.7-8a}$$

$$\;=\; x_1\, c_1\, m_2 + x_2\, c_2\, m_1 \tag{37.7-8b}$$

$$\;=\; x_1\,(1 - c_2\, m_1) + x_2\, c_2\, m_1 \tag{37.7-8c}$$

$$\;=\; x_1 + (x_2 - x_1)\,(m_1^{-1} \mod m_2)\, m_1 \tag{37.7-8d}$$

The last equality is used in the code.

37.8 Quadratic residues

Let p be a prime. The *quadratic residues modulo p* are those values a so that the equation

$$x^2 \equiv a \pmod{p} \tag{37.8-1}$$

has a solution. If the equation has no solution then a is called a *quadratic non-residue modulo p*. A quadratic residue is a square (modulo p) of some number, so we can safely just call it a *square modulo p*.

Let g be a primitive root (the particular choice does not matter), then every nonzero element x can uniquely be written as $x = g^e$ where $0 < e < p$. Rewriting equation 37.8-1 as $x^2 = (g^e)^2 = g^{2e} = a$ it is apparent that the quadratic residues are the even powers of g. The non-residues are the odd powers of g. All generators are non-residues: $g = g^1$.

Let us compute $f(x) := x^{(p-1)/2}$ for both residues and non-residues: With a quadratic residue g^{2e} we get $f(g^{2e}) = g^{2e(p-1)/2} = 1^e = 1$ where we used $g^{p-1} = 1$. With a non-residue $a = g^k$, k odd, we get $f(a) = f(g^k) = g^{k(p-1)/2} = -1$ where we used $g^{(p-1)/2} = -1$ (the only square root of 1 apart from 1 is -1) and $-1^k = -1$ for k odd.

Apparently we just found a function that can tell residues from non-residues. In fact, we rediscovered the so-called *Legendre symbol* usually written as $\left(\frac{a}{p}\right)$. A surprising result (proved by Gauss) about the Legendre symbol is the *law of quadratic reciprocity*: let p and q be distinct odd primes, then

$$\left(\frac{p}{q}\right) = (-1)^{\frac{p-1}{2}\,\frac{q-1}{2}} \left(\frac{q}{p}\right) \tag{37.8-2}$$

Moreover,

$$\left(\frac{-1}{p}\right) = (-1)^{\frac{p-1}{2}} = \begin{cases} +1 & \text{if} \quad p \equiv 1 \pmod 4 \\ -1 & \text{if} \quad p \equiv 3 \pmod 4 \end{cases} \tag{37.8-3}$$

and

$$\left(\frac{2}{p}\right) = (-1)^{\frac{p^2-1}{8}} = \begin{cases} +1 & \text{if} \quad p \equiv \pm 1 \pmod 8 \\ -1 & \text{if} \quad p \equiv \pm 3 \pmod 8 \end{cases} \tag{37.8-4}$$

Further,

$$\left(\frac{3}{p}\right) = 1 \quad \Longleftrightarrow \quad p \equiv \pm 1 \bmod 12 \tag{37.8-5}$$

and

$$\left(\frac{-3}{p}\right) = 1 \quad \Longleftrightarrow \quad p = 2,\ p = 3,\ \text{or}\ p \equiv 1 \bmod 3 \tag{37.8-6}$$

If a is a square modulo p then the polynomial $x^2 - a$ factors as $(x - r_1)(x - r_2)$ where $r_1^2 \equiv a$ and $r_2^2 \equiv a$. Modulo $41 = 4 \cdot 10 + 1$ minus one is a square and we have $x^2 + 1 = (x - 9)(x - 32)$. The polynomial $x^2 + 1$ with coefficients modulo $43 = 4 \cdot 10 + 3$ is irreducible.

The relation between the Legendre symbols of positive and negative arguments is

$$\left(\frac{-a}{p}\right) = (-1)^{\frac{p-1}{2}} \left(\frac{a}{p}\right) = \begin{cases} +\left(\frac{a}{p}\right) & \text{if} \quad p = 4k + 1 \\ -\left(\frac{a}{p}\right) & \text{if} \quad p = 4k + 3 \end{cases} \tag{37.8-7}$$

Modulo a prime $p = 4\,k + 3$, if $+a$ is a square then $-a$ is not a square. The orders of any two elements $+a$ and $-a$ differ by a factor of two. Non-residues can easily be found: $-(b^2)$ is a non-residue for all b.

Modulo a prime $p = 4\,k + 1$, if $+a$ is a square then $-a$ is also a square. The orders of two non-residues $+a$ and $-a$ are identical. The orders of two residues $+a$ and $-a$ can be identical or differ by a factor of two.

A special case are primes of the form $p = 2^x + 1$, the *Fermat primes*. Only five Fermat primes are known today: $2^1 + 1 = 3$, $2^2 + 1 = 5$, $2^4 + 1 = 17$, $2^8 + 1 = 257$ and $2^{16} + 1 = 65537$. To be prime the exponent x must be a power of two. The primitive roots are exactly the non-residues: the maximal order equals $R = \varphi(p) = 2^x$. There are $\varphi\left(\varphi(p)\right) = 2^{x-1}$ primitive roots. There are $(p-1)/2 = 2^{x-1}$ squares which all have order at most $R/2$. Remain 2^{x-1} non-residues which must all be primitive roots.

```
 b\a   0   2   4   6   8  10  12  14  16  18  20  22  24  26  28  30  32  34  36

  0:   0 + 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
  1:   + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + +
  2:   0 + 0 - 0 - 0 + 0 + 0 - 0 + 0 + 0 - 0 - 0 + 0 + 0 - 0 - 0 + 0 + 0 - 0
  3:   0 + - 0 + - 0 + - 0 + - 0 + - 0 + - 0 + - 0 + - 0 + - 0 + - 0 + - 0 + - 0
  4:   0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0
  5:   0 + - - + 0 + - - + 0 + - - + 0 + - - + 0 + - - + 0 + - - + 0 + - - + 0 +
  6:   0 + 0 0 0 + 0 + 0 0 0 + 0 - 0 0 0 - 0 - 0 0 0 - 0 + 0 0 0 + 0 + 0 0 0 + 0
  7:   0 + + - + - - 0 + + - + - - 0 + + - + - - 0 + + - + - - 0 + + - + - - 0 +
  8:   0 + 0 - 0 - 0 + 0 + 0 - 0 - 0 + 0 + 0 - 0 - 0 + 0 + 0 - 0 - 0 + 0 + 0 - 0
  9:   0 + + 0 + + 0 + + 0 + + 0 + + 0 + + 0 + + 0 + + 0 + + 0 + + 0 + + 0 + + 0
 10:   0 + 0 + 0 0 0 - 0 + 0 - 0 + 0 0 0 - 0 - 0 - 0 - 0 0 0 + 0 - 0 + 0 - 0 0 0
 11:   0 + - + + + - - - + - 0 + - + + + - - - + - 0 + - + + + - - - + - 0 + - +
 12:   0 + 0 0 0 - 0 + 0 0 0 - 0 + 0 0 0 - 0 + 0 0 0 - 0 + 0 0 0 - 0 + 0 0 0 - 0
 13:   0 + - + + - - - - + + - + 0 + - + + - - - - + + - + 0 + - + + - - - - + +
 14:   0 + 0 + 0 + 0 0 0 + 0 - 0 + 0 + 0 - 0 + 0 0 0 + 0 + 0 + 0 - 0 - 0 - 0 0 0
 15:   0 + + 0 + 0 0 - + 0 0 - 0 - - 0 + + 0 + 0 0 - + 0 0 - 0 - - 0 + + 0 + 0 0
 16:   0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0
 17:   0 + + - + - - - + + - - - + - + + 0 + + - + - - - + + - - - + - + + 0 + +
 18:   0 + 0 0 0 - 0 + 0 0 0 - 0 - 0 0 0 + 0 - 0 0 0 + 0 + 0 0 0 - 0 + 0 0 0 - 0
 19:   0 + - - + + + + - + - + - - - - + + - 0 + - - + + + + - + - + - - - - + +
 20:   0 + 0 - 0 0 0 - 0 + 0 + 0 - 0 0 0 - 0 + 0 + 0 - 0 0 0 - 0 + 0 + 0 - 0 0 0
```

Figure 37.8-A: Kronecker symbols $\left(\frac{a}{b}\right)$ for small positive a and b.

We will not pursue the issue, but it should be noted that there are more efficient ways than powering to compute f and the concept of quadratic residues can be carried over to composite moduli: the so-called *Kronecker symbol* generalizes the Legendre symbol. An efficient implementation for its computation (following [90, p.29]) is given in [FXT: `kronecker()` in mod/kronecker.cc]:

```
1   int
2   kronecker(umod_t a, umod_t b)
3   // Return Kronecker symbol (a/b).
4   // Equal to Legendre symbol (a/b) if b is an odd prime.
5   {
6       static const int  tab2[] = {0, 1, 0, -1, 0, -1, 0, 1};
7       // tab2[ a & 7 ] := (-1)^((a^2-1)/8)
8
9       if ( 0==b )  return (1==a);
10      if ( 0==((a|b)&1) )  return 0;  // a and b both even ?
11
12      int v = 0;
13      while ( 0==(b&1) )  { ++v;  b>>=1; }
14
15      int k;
16      if ( 0==(v&1) )  k = 1;
17      else             k = tab2[ a & 7 ];
18
19      while ( 1 )
20      {
21          if ( 0==a )  return ( b>1 ? 0 : k );
22
23          v = 0;
24          while ( 0==(a&1) )  { ++v;  a>>=1; }
25
26          if ( 1==(v&1) )  k *= tab2[ b & 7 ];  // k *= (-1)**((b*b-1)/8)
27
28          if ( a & b & 2 )  k = -k;  // k = k*(-1)**((a-1)*(b-1)/4)
29
30          umod_t r = a; // signed:  r = abs(a)
31          a = b % r;
32          b = r;
33      }
34  }
```

A table of Kronecker symbols $\left(\frac{a}{b}\right)$ for small a and b is shown in figure 37.8-A. It was created with the program [FXT: mod/kronecker-demo.cc].

Whether a given number a is a square modulo 2^x can be determined via the simple routine [FXT: `is_quadratic_residue_2ex()` in mod/quadresidue.cc]:

```
1  bool is_quadratic_residue_2ex(umod_t a, ulong x)
2  // Return whether a is quadratic residue mod 2**x
3  {
4      if ( x==1 )  return true;
5      if ( (x>=3 ) && (1==(a&7)) )  return true;
6      if ( (x==2 ) && (1==(a&3)) )  return true;
7      return false;
8  }
```

A curious observation regarding quadratic residues is that exactly for the 29 moduli

```
2, 3, 4, 5, 8, 12, 15, 16, 24, 28, 40, 48, 56, 60, 72, 88, 112, 120,
168, 232, 240, 280, 312, 408, 520, 760, 840, 1320, 1848
```

all quadratic residues are non-prime. This sequence is entry A065428 of [240]. It can be generated using the program [FXT: mod/mod-residues-demo.cc].

See any textbook on number theory for the details of the theory of quadratic residues, [90], [174], and [239] for the corresponding algorithms. A method for watermarking that uses quadratic residues is discussed in [21]. An algorithm to compute conference matrices via quadratic residues is given in section 18.2 on page 367.

37.9 Computation of a square root modulo m

37.9.1 Prime modulus

The square root of a square a modulo a prime $p = 4k + 3$ can be computed as

$$\sqrt{a} \quad = \quad \pm a^{(p+1)/4} \tag{37.9-1}$$

Write $(a^{(p+1)/4})^2 = a^{(p+1)/2} = a^{(p-1)/2+1} = \pm 1 \cdot a = \pm a$, if a is not a square then the square root of $-a$ is obtained. Similar expressions for square root modulo p are developed in [3]. An algorithm for the computation of a square root modulo a prime p (without restriction on the form of p) is given in [90, p.32]. We just give a C++ implementation [FXT: `sqrt_modp()` in mod/sqrtmod.cc]:

```
1  umod_t
2  sqrt_modp(umod_t a, umod_t p)
3  // Return x such that x*x==a (mod p)
4  // p must be an odd prime.
5  // If a is not a square mod p then return 0.
6  {
7      if ( 1!=kronecker(a,p) )  return 0;  // not a square mod p
8
9      // initialize q,t so that  p == q * 2^t + 1
10     umod_t q;  int t;
11     n2qt(p, q, t);
12
13     umod_t z = 0,  n = 0;
14     for (n=1; n<p; ++n)
15     {
16         if ( -1==kronecker(n, p) )
17         {
18             z = pow_mod(n, q, p);
19             break;
20         }
21     }
22
23     if ( n>=p )  return 0;
24
25     umod_t y = z;
```

```
26      uint r = t;
27      umod_t x = pow_mod(a, (q-1)/2, p);
28      umod_t b = x;
29      x = mul_mod(x, a, p);
30      b = mul_mod(b, x, p);
31
32      while ( 1 )
33      {
34          if ( 1==b )  return x;
35
36          uint m;
37          for (m=1; m<r; ++m)
38          {
39              if ( 1==pow_mod(b, 1ULL<<m, p) )  break;
40          }
41
42          if ( m==r )  return  0;  // a is not a square mod p
43
44          umod_t v = pow_mod(y, 1ULL<<(r-m-1), p);
45          y = sqr_mod(v, p);
46          r = m;
47          x = mul_mod(x, v, p);
48          b = mul_mod(b, y, p);
49      }
50  }
```

37.9.2 Prime power modulus

For the computation of a square root modulo a prime power p^x the Newton iteration can be used (see section 28.1.5 on page 561). The case $p = 2$ has to be treated separately [FXT: `sqrt_modpp()` in mod/sqrtmod.cc]:

```
1   umod_t
2   sqrt_modpp(umod_t a, umod_t p, long ex)
3   // Return r so that r^2 == a (mod p^ex)
4   // return 0 if there is no such r
5   {
6       umod_t r;
7
8       if ( 2==p )  // case p==2
9       {
10          if ( false==is_quadratic_residue_2ex(a, ex) )  return 0;  // no sqrt exists
11          else  r = 1;  // (1/r)^2  = a mod 2
12      }
13      else        // case p odd
14      {
15          umod_t z = a % p;
16          r = sqrt_modp(z, p);
17          if ( r==0 )  return  0;  // no sqrt exists
18      }
19      // here r^2 == a (mod p)
20
21      if ( 1==ex )  return  r;
```

Here r is a square root of a modulo p, Newton steps are used to compute $\sqrt{a}$ modulo powers of p:

```
1       const umod_t m = ipow(p, ex);
2       if ( 2==p )  // case p==2
3       {
4           long x = 1;
5           while ( x<ex )  // Newton iteration for inverse sqrt, 2-adic case
6           {
7               umod_t z = a;
8               z = mul_mod(z, r, m);  // a*r
9               z = mul_mod(z, r, m);  // a*r*r
10              z = sub_mod(3, z, m);  // 3 - a*r*r
11              r = mul_mod(r, z/2, m);  // r*(3 - a*r*r)/2 = r*(1 + (1-a*r*r)/2)
12              x *= 2;  //  (1/r)^2 == a mod 2^x
13          }
14          r = mul_mod(r, a, m);
15      }
16      else   // case p odd
17      {
```

```
18          const umod_t h = inv_modpp(2, p, ex);  // 1/2
19          long x = 1;
20          while ( x<ex )  // Newton iteration for square root
21          {
22              umod_t ri = inv_modpp(r, p, ex);    // 1/r
23              umod_t ar = mul_mod(a, ri, m);      // a/r
24              r = add_mod(r, ar, m);              // r+a/r
25              r = mul_mod(r, h, m);               // (r+a/r)/2
26              x *= 2;  //  r^2 == a mod p^x
27          }
28      }
29      return  r;
30  }
```

37.9.3 Arbitrary modulus

The square root modulo an arbitrary number can be computed from the square roots of its prime power factors using the Chinese remainder theorem (see section 37.7 on page 769) [FXT: `sqrt_modf()` in mod/sqrtmod.cc]:

```
1   umod_t
2   sqrt_modf(umod_t a, const factorization &mf)
3   // Return sqrt(a) mod m, given the factorization mf of m
4   {
5       ALLOCA(umod_t, x, mf.nprimes() );
6       for (int i=0; i<mf.nprimes(); ++i)
7       {
8           // x[i]=sqrt(a) modulo i-th prime power:
9           x[i] = sqrt_modpp( a, mf.prime(i), mf.exponent(i) );
10          if ( x[i]==0 )  return 0;  // no sqrt exists
11      }
12      return  chinese(x, mf);  // combine via CRT
13  }
```

37.10 The Rabin-Miller test for compositeness

We describe a probabilistic method to prove compositeness of an integer.

37.10.1 Pseudoprimes and strong pseudoprimes

For a prime p the maximal order of an element equals $p-1$. That is, for all $a \neq 0$

$$a^{p-1} \equiv 1 \mod p \tag{37.10-1}$$

If for a given number n one succeeds to find an $a > 1$ so that $a^{n-1} \neq 1 \bmod p$ the compositeness of n has been proven. Composite numbers n for which $a^{n-1} = 1 \bmod n$ are called *pseudoprime to base a* (or a-pseudoprime). For example, for $n = 15$ one finds

a :	2	3	4	5	6	7	8	9	10	11	12	13	14
a^{14}:	4	9	1	10	6	4	4	6	10	1	9	4	1

We found that 15 is pseudoprime to the bases 4, 11 and 14 which we also could have read off the rightmost column of figure 37.6-B on page 765.

The bad news is that some composite numbers are pseudoprime to very many bases. The smallest such example is number 561 which is pseudoprime to all bases a with $\gcd(a, n) = 1$. Numbers with that property are called *Carmichael numbers*. The first few are 561, 1105, 1729, 2465, 2821, 6601, 8911, ..., this is sequence A002997 of [240]. There are infinitely many Carmichael numbers as proved in [7]. Finding a base that proves a Carmichael number composite is as difficult as finding a factor.

A significantly better algorithm can be found by a rather simple variation. Write $n-1 =: q \cdot 2^t$ where q is odd, we examine the sequence $b := a^q,\ b^2,\ b^4,\ \ldots,\ b^{2^{t-1}} = a^{(n-1)/2}$. We say that n is a *strong pseudoprime to base a* if either $b \equiv 1$ or $b^{2^e} \equiv -1 \equiv n-1$ for some e where $0 \le e < t$. We abbreviate strong pseudoprime as SPP. If neither of the conditions holds then n is proven composite. Then n is either not a pseudoprime to base a or we found a square root of 1 that is not equal to $n-1$.

With two different square roots s_1, s_2 modulo n of some number z (here $z = 1$) one has

$$s_1^2 - z \equiv 0 \mod n \tag{37.10-2a}$$
$$s_2^2 - z \equiv 0 \mod n \tag{37.10-2b}$$
$$s_1^2 - s_2^2 = (s_1 + s_2)\,(s_1 - s_2) \equiv 0 \mod n \tag{37.10-2c}$$

So both $s_1 + s_2$ and $s_1 - s_2$ are nontrivial factors of n if $s_1 \neq n - s_2$. Thereby a square root $s \neq -1$ of 1 proves compositeness because both $\gcd(s+1, n)$ and $\gcd(s-1, n)$ are nontrivial factors of n.

Let $B = \left[b, b^2, b^4, \ldots, b^{2^t}\right]$, then for n prime the sequence B must have one of the following forms: either

$$B = [1, 1, 1, \ldots, 1] \quad \text{or} \tag{37.10-3a}$$
$$B = [*, \ldots, *, -1, 1, \ldots, 1] \tag{37.10-3b}$$

where an asterisk denotes any number not equal to $\pm 1 \bmod n$ (notation as in [174]). For n composite the sequence B can also be of the form

$$B = [*, \ldots, *] \quad (a^{n-1} \neq 1) \quad \text{or} \tag{37.10-4a}$$
$$B = [*, \ldots, *, 1, \ldots, 1] \quad \text{(found square root of 1 not equal to } -1) \tag{37.10-4b}$$

If one of the latter two forms is encountered then n must be composite.

With our example $n = 15$ we have $n - 1 = 7 \cdot 2^1$, thereby $q = 7$ and $t = 1$. We only have to examine the value of b. Values of a for which b is not equal to either $+1$ or -1 prove the compositeness of 15.

a:	2	3	4	5	6	7	8	9	10	11	12	13	14
b:	8	12	4	5	6	13	2	9	10	11	3	7	−1

In our example all bases $\neq 14$ prove 15 composite. As n is always a SPP to base $a = n - 1 \equiv -1$ we restrict our attention to values $2 \le a \le n - 2$.

A pari/gp implementation of the test whether n is a SPP to base a:

```
1   sppq(n, a)=
2   { /* Return whether n is a strong pseudoprime to base a */
3       local(q, t, b, e);
4       q = n-1;
5       t = 0;
6       while ( 0==bitand(q,1), q/=2; t+=1 );
7       /* here  n==2^t*q+1 */
8
9       b = Mod(a, n)^q;
10      if ( 1==b, return(1) );
11      e = 1;
12      while ( e<t,
13          if( (b==1) || (b==n-1), break(); );
14          b *= b;
15          e++;
16      );
17      return( if ( b!=(n-1), 0, 1 ) );
18  }
```

The Carmichael number 561 ($561 - 1 = 35 \cdot 2^4$, so $q = 35$ and $t = 4$) is a SPP to only 8 out of the 558 interesting bases, and not a SPP for any $2 \le a \le 20$ as shown in figure 37.10-A. Note that with $a = 4$ we found $s = 67$ where $s^2 \equiv 1 \bmod 561$ and thereby the factors $\gcd(67+1, 561) = 17$ and $\gcd(67-1, 561) = 33$ of 561.

```
    a=2:      b=263  b^2=166  b^4= 67
    a=3:      b= 78  b^2=474  b^4=276
    a=4:      b=166  b^2= 67  b^4=  1          all SPP basis:
    a=5:      b= 23  b^2=529  b^4=463        a=50:     b=560
    a=6:      b=318  b^2=144  b^4=540        a=101:    b=560
    a=7:      b=241  b^2=298  b^4=166        a=103:    b=  1
    a=8:      b=461  b^2=463  b^4= 67        a=256:    b=  1
    a=9:      b=474  b^2=276  b^4=441        a=305:    b=560
    a=10:     b=439  b^2=298  b^4=166        a=458:    b=560
    a=11:     b=209  b^2=484  b^4=319        a=460:    b=  1
    a=12:     b= 45  b^2=342  b^4=276        a=511:    b=  1
    a=13:     b=208  b^2= 67  b^4=  1
    a=14:     b=551  b^2=100  b^4=463
    a=15:     b=111  b^2=540  b^4=441
    a=16:     b= 67  b^2=  1
    a=17:     b=527  b^2= 34  b^4= 34
    a=18:     b=120  b^2=375  b^4=375
    a=19:     b= 76  b^2=166  b^4= 67
    a=20:     b=452  b^2=100  b^4=463
```

Figure 37.10-A: The Carmichael number $561 = 35 \cdot 2^4 + 1$ is a strong pseudoprime to 8 out of 558 bases a (right) and no basis $2 \le a \le 20$ (left).

37.10.2 The Rabin-Miller test

The *Rabin-Miller test* is an algorithm to prove compositeness of a number n by testing strong pseudoprimality with several bases:

```
1   rm(n, na=20)=
2   { /* Rabin Miller test */
3       local(a);
4       for (a=2, na+2,
5           if ( a>n-2,  break() );
6           if ( 0==sppq(n, a), return(0) ); /* proven composite */
7       );
8       return(1); /* composite with probability less than 0.25^na */
9   }
```

It can be shown that for a composite number the probability to be a SPP to a 'random' base is at most 1/4. Thereby the compositeness of number can in practice quickly be proven. While the algorithm does *not* prove primality, it can be used to rule out compositeness with a very high probability.

```
       Bases tested:    2   3  5  6  7 10 11 12 13 14 15 17
             91:  [3]                  10    12          17
            133:  [2]                     11 12
            145:  [2]                        12          17
            276:  [2]                     11    13
            286:  [2]       3                             17
            703:  [2]       3        7
            742:  [2]                                 15 17
            781:  [2]          5                          17
            946:  [2]                7                 15
           1111:  [2]             6                       17
           1729:  [2]                   10    12
           2047:  [2]    2                 11
           2806:  [2]          5                 13
           2821:  [2]                        12          17
           3277:  [3]    2                             14 15
           4033:  [2]    2                                17
           4187:  [2]                   10                17
           5662:  [2]          5                          17
           5713:  [2]             6                    14
           6533:  [2]             6    10
           6541:  [2]                                   14 15
           7171:  [2]                                   14    17
           8401:  [2]       3           10
           8911:  [3]       3                    12 13
           9073:  [2]                            12    14
```

Figure 37.10-B: All numbers $\le 10,000$ that are strong pseudoprimes to more than one base $a \le 17$ (omitting bases a that are prime powers).

A list (created with the program [FXT: mod/rabinmiller-demo.cc]) of composites $n \le 10,000$ that are

```
    1,373,653  ==  1 + 2^2 * 3^3 * 7 * 23 * 79
               ==  829 * 1657  == (1 + 2^2 *3^2 *23) * (1 + 2^3 *3^2 *23)

    1,530,787  ==  1 + 2 * 3 * 103 * 2477
               ==  619 * 2473  == (1 + 2 *3 *103)    * (1 + 2^3 *3 *103)

    1,987,021  ==  1 + 2^2 * 3^2 * 5 * 7 * 19 * 83
               ==  997 * 1993  == (1 + 2^2 *3 *83)   * (1 + 2^3 *3 *83)

    2,284,453  ==  1 + 2^2 * 3^2 * 23 * 31 * 89
               == 1069 * 2137  == (1 + 2^2 *3 *89)   * (1 + 2^3 *3 *89)

    3,116,107  ==  1 + 2 * 3^2 * 7^2 * 3533
               ==  883 * 3529  == (1 + 2 *3^2 *7^2)  * (1 + 2^3 *3^2 *7^2)

    5,173,601  ==  1 + 2^5 * 5^2 * 29 * 223
               ==  929 * 5569  == (1 + 2^5 *29)      * (1 + 2^6 *3 *29)

    6,787,327  ==  1 + 2 * 3 * 7 * 13 * 31 * 401
               == 1303 * 5209  == (1 + 2 *3 *7 *31)  * (1 + 2^3 *3 *7 *31)

   11,541,307  ==  1 + 2 * 3 * 7 * 283 * 971
               == 1699 * 6793  == (1 + 2 *3 *283)    * (1 + 2^3 *3 *283)

   13,694,761  ==  1 + 2^3 * 3^2 * 5 * 109 * 349
               == 2617 * 5233  == (1 + 2^3 *3 *109)  * (1 + 2^4 *3 *109)
```

Figure 37.10-C: The first composite numbers that are SPP to both bases 2 and 3.

```
      25,326,001  ==  1 + 2^4 * 3^3 * 5^3 * 7 * 67
                  == 2251 * 11251  == (1 + 2 *3^2 *5^3)  * (1 + 2 *3^2 *5^4)

     161,304,001  ==  1 + 2^6 * 3 * 5^3 * 11 * 13 * 47
                  == 7333 * 21997  == (1 + 2^2 *3 *13 *47)  * (1 + 2^2 *3^2 *13 *47)

     960,946,321  ==  1 + 2^4 * 3 * 5 * 29 * 101 * 1367
                  == 11717 * 82013  == (1 + 2^2 *29 *101)  * (1 + 2^2 *7 *29 *101)

   1,157,839,381  ==  1 + 2^2 * 3^3 * 5 * 401 * 5347
                  == 24061 * 48121  == (1 + 2^2 *3 *5 *401)  * (1 + 2^3 *3 *5 *401)

   3,215,031,751  ==  1 + 2 * 3^4 * 5^3 * 7 * 37 * 613
                  == 151 * 751 * 28351
                  == (1 + 2 *3 *5^2)  * (1 + 2 *3 *5^3)  * (1 + 2 *3^4 *5^2 *7)

   3,697,278,427  ==  1 + 2 * 3^3 * 31 * 563 * 3923
                  == 30403 * 121609  == (1 + 2 *3^3 *563)  * (1 + 2^3 *3^3 *563)
```

Figure 37.10-D: All composite numbers $n \leq 2^{32}$ that are SPP to the three bases 2, 3 and 5.

SPP to more than one base $a \leq 17$ is shown in figure 37.10-B. The table indicates how effective the Rabin-Miller algorithm actually is: it does not contain a single number pseudoprime to both 2 and 3. The first few odd composite numbers that are SPP to both bases $a = 2$ and $a = 3$ are shown in figure 37.10-C. There are 104 such composite $n < 2^{32}$, given in [FXT: data/pseudo-spp23.txt]. This sequence of numbers is entry A072276 of [240], entry A001262 gives the base-2 SPPs, and entry A020229 gives the base-3 SPPs. We note the uneven distribution modulo 12:

```
(n%12: num)     (1: 75)     (5: 9)     (7: 18)     (11: 2)
```

Composites that are SPP to the three bases 2, 3 and 5 are quite rare, figure 37.10-D shows all 6 such composite numbers smaller than 2^{32} (values taken from [213] which lists all such numbers $< 25 \cdot 10^9$). Thereby one can speed up the Rabin-Miller test for small values of n (say, $n < 2^{32}$) by only testing the bases $a = 2, 3, 5$ and, if n is a SPP to these bases, look up the composites in the table. The smallest odd composites that are SPP to the first k prime bases up to $k = 8$ are determined in [152], they are given as sequence A006945 of [240].

```
     composite       SPP to base
          2047       2
       1373653       2, 3
      25326001       2, 3, 5
    3215031751       2, 3, 5, 7
 2152302898747       2, 3, 5, 7, 11
 3474749660383       2, 3, 5, 7, 11, 13
341550071728321      2, 3, 5, 7, 11, 13, 17 [and 19]
341550071728321      2, 3, 5, 7, 11, 13, 17, 19
```

Note that if the probability of a base not proving compositeness was exactly 1/4 we would find *much* more entries in figure 37.10-D. Slightly overestimating the number of composites below N as N, there

should be about $(1/4)^3\,N = N/64$ entries, that is $2^{26} \approx 6 \cdot 10^7$ for $N = 2^{32}$, but we have only six entries. Thereby the Rabin-Miller test is in practice significantly more efficient than one may initially assume. Let $p_{k,t}$ be the probability that a k-bit composite 'survives' t passes of the Rabin-Miller test. Then we have, as shown in [101],

$$p_{k,1} \quad < \quad k^2\, 4^{2-\sqrt{k}} \qquad \text{for} \quad k > 2 \tag{37.10-5}$$

For large numbers, the bound on the left hand side is much smaller than 1/4: for example, $p_{1000,1} < 2^{-39}$. Other bounds given in the cited paper are

$$p_{100,10} \quad < \quad 2^{-44} \tag{37.10-6a}$$

$$p_{300,5} \quad < \quad 2^{-60} \tag{37.10-6b}$$

$$p_{600,1} \quad < \quad 2^{-75} \tag{37.10-6c}$$

The last bound is stronger than that of relation 37.10-5. Still stronger bounds are given in [75], especially the relation $p_{k,t} < 4^{-t}$ for all $k \geq 2$ and $t \geq 1$.

```
      Bases tested:  2   3  5  6  7 10 11 12 13 14 15 17
             11476:      3  5                       15 17
             88831:               7    11 12        15 17
            188191:      3         7 10                17
            597871:                  10 11 12 13
            736291:                  7 10    12 13
            765703:                  10 11 12    14
           1024651:      3  5        7          12       15
           1056331:                  7 10       12 13
           1152271:      3           7     11 12 13
           1314631:         5        7     11 12 13
           1373653:   2  3      6                12             17
           1530787:   2  3      6                12
           1627921:      3  5                             14 15
           1857241:         5                 11 12       14
           1987021:   2  3      6                12             17
           2030341:                           11 12       14 15
           2284453:   2  3      6  7           11 12
           2741311:                           11 12       14 15
           3116107:   2  3      6                12
           4181921:   2     5          10           13
           4224533:               6        11           14 15 17  <--=
           5122133:               6  7 10 11
           5173601:   2  3      6                12
           5481451:      3  5                    12          15
           6594901:         5                    12 13 14
           6787327:   2  3      6                12 13
           8086231:         5                    12    14    17
           9504191:                           11 12 13       15
           9863461:   2     5  6     10                        17
```

Figure 37.10-E: Composites $\leq 10^7$ that are SPP to at least four bases.

The composites $\leq 10^7$ that are SPP to four or more bases $a < 17$ are shown in figure 37.10-E. We omit values of a that are perfect powers because if n is a base-a SPP then it is also a base-a^k SPP for all $k > 1$. The entry for $n = 4224533$ (marked with an arrow) shows that a number that is not a SPP to two bases a_1 and a_2 may still be a SPP to the base $a_1 \cdot a_2$ (here $a_1 = 2$, $a_2 = 3$). This indicates that one might want to restrict the tested bases to primes. All odd composite numbers $\leq 10^7$ that are SPP to four or more prime bases $a \leq 17$ are

```
 Bases tested:   2  3  5  7 11 13 17
1152271:  [4]       3     7 11 13
1314631:  [4]          5  7 11 13
2284453:  [4]    2  3     7 11
```

Note that a number that is a SPP to bases a_1 and a_2 is *not* necessarily SPP to the base $a_1 \cdot a_2$. An example is $n = 9,006,401$ which is a SPP to bases 2 and 5 but not to base 10:

```
9006401:   2     4  5        8                       16    18
```

All composites $\leq 10^7$ that are SPP to bases 2 and 3 are also SPP to base 6, same for bases 2 and 5. Out of six composites $\leq 10^7$ that are SPP to bases 2 and 7 three are not SPP to base 14:

```
  314821:   2     4     6  7  8  9                    16    18
 2269093:   2     4        7  8              14       16
 2284453:   2  3  4     6  7  8  9    11 12           16    18
 3539101:   2     4        7  8                 13     16
 5489641:   2     4        7  8                 14     16
 6386993:   2     4        7  8                 14     16
```

37.10.3 Implementation of the Rabin-Miller test

A C++ implementation of the test for pseudoprimality is given in [FXT: mod/rabinmiller.cc]:

```
bool
is_strong_pseudo_prime(const umod_t n, const umod_t a, const umod_t q, const int t)
// Return whether n is a strong pseudoprime to base a.
// q and t must be set so that  n == q * 2^t + 1
{
    umod_t b = pow_mod(a, q, n);

    if ( 1==b )  return true;  // passed
    // if ( n-1==b )  return true;  // passed

    int e = 1;
    while ( (b!=1) && (b!=(n-1)) && (e<t) )
    {
        b = mul_mod(b, b, n);
        e++;
    }

    if ( b!=(n-1) )  return false;  // =--> composite

    return  true;  // passed
}
```

It uses the routine

```
void
n2qt(const umod_t n, umod_t &q, int &t)
// Set q,t so that  n == q * 2^t + 1
// n must not equal 1, else routine loops.
{
    q = n - 1;  t = 0;
    while ( 0==(q & 1) )  { q >>= 1; ++t; }
}
```

Now the Rabin-Miller test can be implemented as

```
bool
rabin_miller(umod_t n, uint cm/*=0*/)
// Rabin-Miller compositeness test.
// Return true of none of the bases <=cm prove compositeness.
// If false is returned then n is proven composite (also for n=1 or n=0).
// If true is returned the probability
//   that n is composite is less than (1/4)^cm
{
    if ( n<=1 )  return false;
    if ( n < small_prime_limit )  return  is_small_prime( (ulong)n );

    umod_t q;
    int t;
    n2qt(n, q, t);

    if ( 0==cm )  cm = 20;  // default
    uint c = 0;
    while ( ++c<=cm )
    {
        umod_t a = c + 1;

        // if n is a c-SPP then it also is a c**k (k>1) SPP.
        // That is, powers of a non-witness are non-witnesses.
        // So we skip perfect powers:
        if ( is_small_perfpow(a) )  continue;

        if ( a >= n )  return  true;
        if ( !is_strong_pseudo_prime(n, a, q, t) )  return false;  // proven composite
    }

    return true;  // strong pseudoprime for all tested bases
```

```
32  }
```

The function `is_small_perfpow()` [FXT: mod/perfpow.cc] returns `true` if its argument is a (small) perfect power. It uses a lookup in a precomputed bit-array.

A generalization of the Rabin-Miller test applicable when more factors (apart from 2) of $n-1$ are known is given in [40]. Another generalization (named extended quadratic Frobenius primality test) is suggested in [102].

37.11 Proving primality

We describe several methods to prove primality. Only the first, Pratt's certificate of primality, is applicable for numbers of arbitrary form but not practical in general because it relies on the factorization of $n-1$. The Pocklington-Lehmer test only needs a partial factorization of $n-1$. We give further tests applicable only for numbers of a special forms: Pepin's test, the Lucas-Lehmer test, and the Lucas test.

As already said, the Rabin-Miller test can only prove compositeness. Even if a candidate 'survives' many passes, we only know that it is prime with a high probability.

37.11.1 Pratt's certificate of primality

Only with a prime modulus p the maximal order equals $R = p-1$. To determine the order of an element modulo p one needs the factorization of $p-1$. Now if the factorization of $p-1$ is known and we can find a primitive root then we do know that p is prime. Thereby it is quite easy to prove primality for numbers of certain special forms. For example, let $p := 2 \cdot 3^{30} + 1 = 411,782,264,189,299$. One finds that 3 is a primitive root and so we know that p is prime.

```
[314159311, [3], [2, 3, 5, 199, 1949]]
    [2, "--"]
    [3, [2], [2]]
        [2, "--"]
    [5, [2], [2]]
        [2, "--"]
    [199, [3], [2, 3, 11]]
        [2, "--"]
        [3, [2], [2]]
            [2, "--"]
        [11, [2], [2, 5]]
            [2, "--"]
            [5, [2], [2]]
                [2, "--"]
    [1949, [2], [2, 487]]
        [2, "--"]
        [487, [3], [2, 3]]
            [2, "--"]
            [3, [2], [2]]
                [2, "--"]
```

Figure 37.11-A: A certificate for the primality of $p = 314,159,311$.

In general, the factorization of $p-1$ can contain large factors whose primality needs to be proven. Recursion leads to a primality certificate in the form of a tree which is called *Pratt's certificate of primality*.

A certificate for the primality of $p = 314,159,311$ is shown in figure 37.11-A. The first line says that 3 is a primitive root of $p = 314159311$ and $p-1$ has the prime factors $2, 3, 5, 199, 1949$ (actually, $p-1 = 2 \cdot 3^4 \cdot 5 \cdot 199 \cdot 1949$ but we can ignore exponents). In the second level (indented by 4 characters) appear the

```
    [3141592653589793238462643383279502884197169399375 31, [3], \
        [2, 3, 5, 67, 89, 151, 39829177707048956693, 29200179492960384562 1939]]
        [151, [6], [2, 3, 5]]
        [39829177707048956693, [2], [2, 9957294426762239173]]
            [9957294426762239173, [6], [2, 3, 7, 11, 14153, 385109, 1977139]]
                [14153, [3], [2, 29, 61]]
                [385109, [2], [2, 43, 2239]]
                    [2239, [3], [2, 3, 373]]
                        [373, [2], [2, 3, 31]]
                [1977139, [3], [2, 3, 109841]]
                    [109841, [3], [2, 5, 1373]]
                        [1373, [2], [2, 7]]
        [29200179492960384562 1939, [2], [2, 13, 3157127, 3557296955910619]]
            [3157127, [7], [2, 7, 225509]]
                [225509, [2], [2, 56377]]
                    [56377, [5], [2, 3, 29]]
            [3557296955910619, [3], [2, 3, 47, 673, 6247908971]]
                [673, [5], [2, 3, 7]]
                [6247908971, [2], [2, 5, 624790897]]
                    [624790897, [5], [2, 3, 13016477]]
                        [13016477, [2], [2, 11, 29, 101]]
                            [101, [2], [2, 5]]
```

Figure 37.11-B: A shortened certificate for the primality of first prime greater than $\pi \cdot 10^{50}$. Here all primes smaller than 100 are considered trivially verifiable and not listed.

prime factors just determined together with their primality certificates: the prime 2 is trivially accepted, all other primes are followed by their (further indented) certificates.

The certificate was produced with the following pari/gp code:

```
1   indprint(x, ind)=
2   { /* print x, indented by ind characters */
3       for (k=1, ind, print1(" ") );
4       print(x);
5   }
6
7   pratt(p, ind=0)=
8   {
9       local( a, p1, f, nf, t );
10      if ( p<=2,   \\ 2 is trivially prime
11          indprint([p, "--"], ind);
12          return();
13      );
14      \\ p-1 is factored here:
15      a = component(znprimroot(p),2);
16      \\ but we cannot access the factorization, so we do it "manually":
17      p1 = p-1;
18      f = factor(p1);
19      nf = matsize(f)[1];
20      t = vector(nf,j, f[j,1] );  f = t;  \\ prime factors only
21      indprint([p, [a], t], ind);
22      \\ recurse on prime factors of p-1:
23      for (k=1, nf, pratt(f[k], ind+4));
24      return();
25  }
```

```
? p=nextprime(Pi*10^8);
? pratt(p)
```

The routine has to be taken with a grain of salt as we rely on the fact that `znprimroot(p)` will fail for composite p:

```
? pratt(1000)
  ***   primitive root does not exist in gener
```

The routine has an additional parameter `ind` determining the indentation used with printing. This parameter is incremented with the recursion level, resulting in the tree-like structure of the output. This

little trick is often useful with recursive procedures.

With a precomputed table of small primes (see section 37.3 on page 759) the line

```
if ( p<=2,   \\ 2 is trivially prime
```

can be changed to something like

```
if ( (p<=ptable_max) && (ptable[p]==1),   \\ trivial to verify
```

which will shorten the certificate significantly. A certificate for the smallest prime p greater than $\pi \cdot 10^{50}$ and `ptable_max`=100 is shown in figure 37.11-B, the output of 'trivial' primes is suppressed. We note that $p = \lceil \pi \cdot 10^{50} \rceil + 20$.

Once a certificate is computed it can be verified very quickly. Note that this type of primality certificate needs the factorization of $p-1$ and thus is usually not feasible for large values of p.

37.11.2 The Pocklington-Lehmer test

Let $p - 1 = F \cdot U$ where $F > U$ and the complete factorization of F is known. If, for each prime factor q of F, we can find a_q such that $a_q^{p-1} \equiv 1 \bmod p$ and $\gcd\left(a_q^{(p-1)/q} - 1, p\right) = 1$, then p is prime.

The following implementation removes entries from the list of prime factors q of F until the list is empty:

```
1   pocklington_lehmer(F, u, c=10000)=
2   { /* Pocklington-Lehmer test for the primality of p=f*u+1.
3      * Return last successful base, else zero.
4      * F must be the factorization of f.
5      * Test bases a=2...c
6      * Must have u<f.
7      */
8       local(n, f, C, p, t, ct);
9       n = matsize(F)[1];
10      f = prod(j=1, n, F[j,1]^F[j,2]);
11      if ( f<=u, return(0) );
12      p = f*u + 1;
13      C = vector(n, j, (p-1)/F[j,1]);
14      ct = n;  \\ number remaining prime divisors of f
15      for (a=2, c,
16          if ( 1==Mod(a,p)^(p-1),
17              for (j=1, n,
18                  if ( C[j]!=0,  \\ skip entries already removed
19                      t = component( Mod(a,p)^C[j], 2);
20                      if( 1==gcd(t-1, p),
21                          C[j] = 0;  \\ remove entry
22                          ct -= 1;   \\ number of remaining entries
23                      );
24                  );
25              );
26              if ( ct==0, return(a) );
27          );
28      );
29      return( 0 );
30  }
```

We search all primes of the form $p = F \cdot U + 1$ where $F = 100!$, $U = F - d$, and d lies in the range $1, \ldots, 1000$. Only candidates that are strong pseudoprimes to both bases 2 and 3 are tested:

```
f=100!;
F=factor(f);
{ for (d=1, 1000,
    u = f - d;
    p = f*u+1;
    if ( sppq(p, 2) && sppq(p,3),
        q2 = pocklington_lehmer(F, u);
        print1(d, ":  ");
        print1("  ", q2);
        print();
    );
 ) }
```

We find five such primes $\approx 8.70978248908948 \cdot 10^{315}$ (in about 10 seconds):

```
    d:   last a
  45:    103
 778:    101
 818:    101
 880:    101
 884:    103
```

The returned value a_q is the one that did lead to the removal of the last entry in `C[]`. The value is smaller with less prime factors of F. Using $F = 2^{500}$ we find primes ($\approx 1.07150860718626 \cdot 10^{301}$) of the form $p = F \cdot U + 1$ where $U = F - d$ and $1 \le d \le 3000$ for the following d and maximal a_q:

```
    d:   last a        d:   last a
  214:    5         1383:    3
  294:    3         1801:   13
 1023:    3         2041:   11
 1114:    5         2481:    3
 1321:   17
```

The search takes about 20 seconds. Discarding candidates that have small prime factors ($p < 1,000$) gives a four-fold speedup. The prime $2^{3340}\,(2^{3340} - 1633) + 1 \approx 7.59225935 \cdot 10^{2010}$ is found within 5 minutes.

A further refinement of the test is given in [90], see also [72].

37.11.3 Tests for $n = k\,2^t + 1$

37.11.3.1 Proth's theorem and Pepin's test

For numbers of the form $p = q \cdot 2^t + 1$ with q odd and $2^t > q$ primality can be proven as follows: If there is an integer a so that $a^{(p-1)/2} \equiv -1$ then p must be prime. This is *Proth's theorem*.

The 'FFT-primes' (see section 25.1 on page 525) are natural candidates for Proth's theorem. For example, with $p := 2^{57} \cdot 29 + 1 = 4,179,340,454,199,820,289$ one finds that $a^{(p-1)/2} \equiv -1$ for $a = 3$, so p must be prime. Note that Proth's theorem is the special case $F = 2^t > k = U$ of the Pocklington-Lehmer test.

Numbers of the form 2^t+1 are composite unless t is a power of two. The candidates are therefore restricted to the so-called *Fermat numbers* $F_n := 2^{2^n} + 1$. Here it suffices to test whether $3^x \equiv -1 \bmod F_n$ where $x = 2^{t-1}$:

```
1  pepin(tx)=
2  {
3      local(t, F, x);
4      t = 2^tx;
5      F = 2^t+1;
6      x = 2^(t-1);
7      return(  (-1==Mod(3,F)^x) );
8  }
```

This test is known as *Pepin's test*. As shown in section 37.8 on page 772 all non-residues are primitive roots modulo prime F_n. Three is just the smallest non-residue.

```
for (tx=1,12, print(tx,"  ", pepin(tx)))
   1  1  \\ F_1 = 5
   2  1  \\ F_2 = 17
   3  1  \\ F_3 = 257
   4  1  \\ F_4 = 65537
   5  0
   6  0
   7  0
```

No Fermat prime greater than $F_4 = 65537$ is known today and all F_n where $5 \le n \le 32$ are known to be composite.

Note that F_{n+1} has (about) twice as many bits as F_n. Further the number of squarings involved in the test $(t-1)$ is (about) doubled. If we underestimate the cost of multiplying N-bit numbers with N operations we get a lower bound for the ratio of the costs of testing F_{n+1} and F_n of four. Assuming the computer power doubles every 18 month and Pepin's test of F_n is just feasible today we'd have to wait three years (36 month) before we can test F_{n+1}. The computation that proved F_{24} composite is described in [97].

37.11.3.2 What to consider before doing the Pepin's test

As $2^t \equiv -1 \bmod F_n = 2^t + 1$ we see that the order of two equals $2\,t = 2^{n+1}$. The same is true for factors of composite Fermat numbers. When searching factors of F_n one only needs to consider candidates of the form $1 + k\,2^{n+2}$. A routine that searches for small factors of f_n can be implemented as:

```
ord2pow2(p)=
\\ Return the base-2 logarithm of the order of two modulo p
\\ Must have: ord(2)==2^k for some k
{
    local(m, rx);
    rx = 0;
    m = Mod(2,p);
    while ( m!=1, m*=m; rx++; );
    return( rx );
}
```

```
ftrialx(n, mm=10^5, brn=0)=
\\ Try to find small factors of the Fermat number F_n=2^(2^n)+1
\\ Try factors 1+ps, 1+2*ps, ... , 1+mm*ps  where ps=2^(n+2)
\\ Stop if brn factors were found (zero: do not stop)
{
    local(p,ps,ttx,fct);
    ps = 2^(n+2);  \\ factors are of the form 1+k*ps
    p = ps+1;      \\ trial factor
    ttx = 2^(n+1); \\ will test whether Mod(2,p)^ttx==1
    fct = 0;  \\ how many factors were found so far
    for (ct=1, mm,
        if ( (Mod(2,p)^(ttx)==1)  \\ order condition
            && ( (rx=ord2pow2(p)) == n+1 ) \\ avoid factors of smaller Fermat numbers
        , /* then */
            print1(n, ":  ");
            print1(p);
            print1("    p-1=",factor(p-1));
            print();
            fct++;
            if ( fct==brn,  break() );
        );
        p += ps;
    );
    return(fct);
}
```

We create a list of small prime factors of F_n for $5 \le n \le 32$ where the search is restricted to factors $f \le 1 + 10^5\,2^{n+2}$ and stopped when a factor was found:

```
for(n=5,32, ftrialx(n, 10^5, 1); );
   5:  641    p-1=[2, 7; 5, 1]
   6:  274177    p-1=[2, 8; 3, 2; 7, 1; 17, 1]
   9:  2424833    p-1=[2, 16; 37, 1]
  10:  45592577    p-1=[2, 12; 11131, 1]
  11:  319489    p-1=[2, 13; 3, 1; 13, 1]
  12:  114689    p-1=[2, 14; 7, 1]
  15:  1214251009    p-1=[2, 21; 3, 1; 193, 1]
  16:  825753601    p-1=[2, 19; 3, 2; 5, 2; 7, 1]
  18:  13631489    p-1=[2, 20; 13, 1]
  19:  70525124609    p-1=[2, 21; 33629, 1]
  23:  167772161    p-1=[2, 25; 5, 1]
  32:  25409026523137    p-1=[2, 34; 3, 1; 17, 1; 29, 1]
```

A list for $5 \le n \le 300$ is given in [FXT: data/small-fermat-factors.txt]. Note that an entry

```
 201:  1245698371909569261600129013982869249475211760780421005925626675 21 \
   p-1=[2, 204; 3, 1; 5, 1; 17, 1; 19, 1]
```

asserts the compositeness of the number F_{201} where Pepin's test is out of reach by far. Indeed, its binary representation could not be stored in all existing computer memory combined: F_{201} is a $\log_2(F_{201}) \approx 2^{201} = 3.2138 \cdot 10^{60}$ -bit number.

The status of the (partial) factorizations of Fermat numbers is given in [162].

37.11.4 Tests for $n = k\,2^t - 1$

37.11.4.1 The Lucas-Lehmer test for Mersenne numbers

Define the sequence H by $H_0 = 1$, $H_1 := 2$ and $H_i = 4\,H_{i-1} - H_{i-2}$. The Mersenne number $n = 2^e - 1$ is prime exactly if $H_{2^{e-2}} \equiv 0 \bmod n$. The first few terms of the sequence H are

```
    k:  0  1  2   3   4   5     6     7     8      9      10      11      12
  H_k:  1  2  7  26  97  362  1351  5042  18817  70226  262087  978122  3650401 ...
```

The numbers H_k can be computed efficiently via the index doubling formula $H_{2k} = 2\,H_k^2 - 1$. Starting with the value $H_1 = 2$ and computing modulo n the implementation is a s simple as

```
1   LL(e)=
2   {
3       local(n, h);
4       n = 2^e-1;
5       h = Mod(2,n);
6       for (k=1, e-2, h=2*h*h-1);
7       return( 0==h );
8   }
```

```
? LL(521)
  1  \\  2^521-1 is prime
? LL(239)
  0  \\  2^239-1 is composite
? LL(9941)
  1 \\ 2^9941-1 is prime
? ##
  ***   last result computed in 4,296 ms.
```

The algorithm is called the *Lucas-Lehmer test*. Note that most sources use the sequence $V = 2, 4, 14, 52, 194, 724, 2702, \ldots$ that satisfies the same recurrence relation. We have $H_k = \frac{1}{2}\,V_k$ (H is *half* of V). The index doubling relation becomes $V_{2k} = V_k^2 - 2$. The sequence of values H_{2^k} starts as

```
2, 7, 97, 18817, 708158977, 1002978273411373057, 2011930833870518011412817828051050497, ...
```

This is entry A002812 of [240]; entry A003010 gives the values V_{2^k}:

```
4, 14, 194, 37634, 1416317954, 2005956546822746114, 4023861667741036022825635656102100994, ...
```

37.11.4.2 What to consider before doing the Lucas-Lehmer test

The exponent e of a Mersenne prime must be prime, else n factors algebraically as

$$2^e - 1 \quad = \quad \prod_{d \backslash e} Y_d(2) \tag{37.11-1}$$

where $Y_k(x)$ is the k-th cyclotomic polynomial (see section 35.1.1 on page 677). For example, with $2^{21} - 1 = 2097151$ the following factors are found:

```
? m=1; fordiv(21,d, y2=subst(polcyclo(d,x),x,2); m*=y2;print(d,": ",y2)); m
  1: 1
  3: 7
  7: 127
  21: 2359
  2097151  \\ == 2^21-1
```

The factors found are not necessarily prime: here 2359 factors further into $7 \cdot 337$. More information on the multiplicative structure of $b^e \pm 1$ can be found in [73].

Before doing the Lucas-Lehmer test one should do a special version of trial division based on the following observation: any factor f of $m = 2^e - 1$ has the property that $2^e \equiv 1 \bmod f$. That is $2^e - 1 \equiv 0 \bmod f$, so $m \equiv 0 \bmod f$ and f divides m. We further use the fact that factors are of the form $2\,e\,k + 1$ and that $f \equiv \pm 1 \bmod 8$. The following routine does not try assert the primality of a candidate factor as this would render the computation considerably slower.

```
1   mers_trial(e, mct=10^7, bnf=0)=
2   \\  try to discover small factors of the Mersenne number 2^e-1
3   \\  e : exponent of the Mersenne number
4   \\  mct : how many factors are tried
```

```
5   \\  pfq : stop with the factor found (zero: do not stop)
6   {
7       local(f, fi, ct, fct);
8       print("exponent e=",e);
9       print("trying up to ", mct, " factors");
10      fi=2*e; \\ factors are of the form 2*e*k+1
11      f=1;
12      ct=0;
13      fct=0;  \\ how many factors where found so far
14      while (ct < mct,
15          f += fi;
16          m8 = bitand(f, 7);  \\ factor modulo 8
17          if ( (1!=m8) && (7!=m8), next(); );  \\ must equal +1 or -1
18          if ( Mod(2, f)^e == Mod(1, f),
19              print(f, "  ", isprime(f));  \\ give factor and tell whether it is prime
20              fct++;
21              if ( fct==bnf , break(); );
22         );
23         ct++;
24      );
25  }
```

For $m = 2^{10007} - 1$ (3013 decimal digits) we find three factors of which all are prime:

```
? e=10007; mers_trial(e,,3);
exponent e=10007
trying up to 10000000 factors
  240169  1
  60282169  1
  136255313  1
? ##
  ***   last result computed in 44 ms.
? ceil((e*log(2.0)/log(10.0)))
  3013 \\ m=2^e-1 has 3,013 decimal digits
```

Sometimes one is lucky with truly huge numbers:

```
? e=2^31-1; mers_trial(e,,1);
exponent e=2147483647
trying up to 10000000 factors
  295257526626031  1
? ##
  ***   last result computed in 583 ms.
? ceil((e*log(2.0)/log(10.0)))
  646456993  \\  m=2^e-1 has 646,456,993 decimal digits
```

Note that we found that $m = 2^e - 1$ is prime if and only if there is no prime $f < m$ where the order of two equals e. A special case is sometimes given as follows: if both $p = 4\,k + 3$ and $q = 2\,p + 1$ are prime then q divides $2^p - 1$ (because the order of 2 modulo q equals p).

By the way, if both $p = 4\,k + 1$ and $q = 2\,p + 1$ are prime then q divides $2^p + 1$ (because the order of 2 modulo q equals $2\,p$ and $2^{2\,p} - 1 = (2^p + 1)\,(2^p - 1)$).

37.11.4.3 Lucas-Lehmer test with floats *

The Binet form (see section 34.1.6 on page 665) of the sequence H_n is

$$H_n \;=\; \frac{1}{2}\left[\left(2+\sqrt{3}\right)^n + \left(2+\sqrt{3}\right)^{-n}\right] \tag{37.11-2}$$

We can rewrite the expression in the form

$$H_n \;=\; \frac{1}{2}\left[\exp(x)^n + \exp(x)^{-n}\right] \;=\; \frac{1}{2}\left[\exp(n\,x) + \exp(-n\,x)\right] \tag{37.11-3}$$

where $x = \log(2 + \sqrt{3})$. The hyperbolic cosine can be defined as

$$\cosh(z) \;=\; \frac{1}{2}\left[\exp(z) + \exp(-z)\right] \tag{37.11-4}$$

and the expression equals H_n for $z = n \log(2+\sqrt{3})$. Now we can give a criterion equivalent to the Lucas-Lehmer condition as follows:

$$\cosh\left(2^{m-2} \log\left(2+\sqrt{3}\right)\right) \equiv 0 \mod M_m \quad \Longrightarrow \quad M_m \text{ is prime} \tag{37.11-5}$$

The relation is computationally useless because the quantity to be computed grows doubly-exponentially with m: the number of digits grows exponentially with m. Already for $m = 17$ the calculation has to be carried out with more than 18,741 decimal digits:

```
? cosh(2^(17-2)*log(2+sqrt(3)))
 1.88889395811398377260975384780566028594658443155 51 E18741
```

The program [hfloat: examples/ex8.cc] does the computations in the obvious (insane) way. Using a precision of 32,768 decimal digits we obtain:

```
 cosh(...)=
  +.18888939581139837726097538478056602859465844315551 \
    [... about 18,000 digits ...]
     ... 55797500398006802841700000000000000 ...
                              ^[decimal point after 7]
    00000000000000000000000000000000000000000000000000 \
    [...]
    00000000000000001549695720446140150427588985400185472*10^18742
                    [nonzero due to numerical imprecision]
```

After rounding and computing the modulus the program declares $M_{17} = 2^{17} - 1$ prime. All this using just 4 MB of memory and computations equivalent to about 35 FFTs of length 1 million, taking about 4 seconds. This is many many million times the work needed by the original (sane) version of the test. Even trial division would have been significantly faster.

The number M_{31} would need a bigger machine as the computations needs a precision of more than 300 million digits:

```
? (2^(31-2)*log(2+sqrt(3)))/log(10)  /* approx decimal digits */
 307062001.46039800926268312190009204
```

Apart from being insane the computation can be used to test high precision floating point libraries.

37.11.4.4 The Lucas test

The Lucas-Lehmer test can be generalized for a less restricted set of candidates. The *Lucas test* can be stated as follows (taken from [222, p.131]):

Let $n = k\,2^t - 1$ were k is odd, $2^t > k$, $n \neq 0 \bmod 3$ and $k \neq 0 \bmod 3$ (so we must have $n \equiv 1 \bmod 3$). Then n is prime if and only if $H_{(n+1)/4} \equiv 0 \bmod n$ where H is as given above.

To turn this into an efficient algorithm use the relation $(n+1)/4 = k\,2^{t-2}$. Firstly, compute H_k as described in section 34.1.1 on page 657:

$$[H_k,\ H_{k+1}] \quad = \quad [H_0,\ H_1] \begin{bmatrix} 0 & -1 \\ 1 & 4 \end{bmatrix}^k \tag{37.11-6}$$

This is a one-liner in pari/gp:

```
? H(k)= return( ([1,2] * [0, -1; 1, 4]^k)[1] );
? for(k=0,10,print(k,": ",H(k),"  = 1/2 * ",2*H(k)))
  0: 1  = 1/2 * 2
  1: 2  = 1/2 * 4
  2: 7  = 1/2 * 14
  3: 26  = 1/2 * 52
  4: 97  = 1/2 * 194      /*  = 2*7^2-1 = (14^2-2)/2 */
  5: 362  = 1/2 * 724
  6: 1351  = 1/2 * 2702   /*  = 2*26^2-1 = (52^2-2)/2 */
  7: 5042  = 1/2 * 10084
  8: 18817  = 1/2 * 37634
  9: 70226  = 1/2 * 140452
  10: 262087  = 1/2 * 524174
```

To compute $H_{k\,2^{t-2}}$ from H_k use ($t-2$ times) the index doubling relation $H_{2k} = 2\,H_k^2 - 1$. The test can be implemented as

```
1   H(k, n)= return( (Mod([1,2],n) * Mod([0,-1; 1,4], n)^k)[1] );
2
3   lucas(k, t)=
4   {
5       local(n, h);
6       /* check preconditions: */
7       if ( 0==bitand(k,1), return(0) );  \\ k must be odd
8       if ( k>=2^t, return(0) );
9       n = k*2^t-1;
10      if ( n%3!=1, return(0) );  \\ gcd(3,k)!=0 && gcd(3,n)!=0
11
12      /* main loop: */
13      h = H(k, n);
14      for (j=1,t-2, h*=h; h+=h; h-=1; );  \\ index doubling
15      return ( 0==h );
16  }
```

Note that the routine returns 'false' even for primes if the preconditions are not met. With $n = 5{\cdot}2^{12}-1 = 20479$ we obtain

```
n=20479  k=5  t=12
 j  H_{k*2^j} modulo n
 0      362
 1    16339
 2    17832
 3     5581
 4    18482
 5     9686
 6     8593
 7     5228
 8     5516
 9     9402
10        0
```

which shows that 20479 is prime. Proving $n = 5 \cdot 2^{1340} - 1$ prime takes about 10 milliseconds. The following code finds the first value $t \geq 2500$ so that $n = 5 \cdot 2^t - 1$ is prime:

```
k=5;  t=2500;  while ( 0==lucas(k,t), t+=1; );  t
```

Within 1 second we obtain the result $t = 2548$.

37.11.4.5 Numbers of the form $n = 24\,j + 7$ and $n = 24\,j + 19$ *

```
          n:   SPP bases a<100,000 (max 5 given)
    1037623:  67191 67192                                  [--snip--]
    2211631:  6333 7260 8160 16793 21219 21580          2946282799:
    4196191:  9104 26498 93477                          3075304399:
    7076623:                                            3145717759:
    9100783:                                            3299597407:
   11418991:  44936                                     3554502799:
   15219559:                                            3554889199:
   21148399:                                            4091977039:
    [--snip--]                                          4207009999:
  829577839:
  887557999:  4899 33982 46674 62180
  961315183:
 1192222639:
  [--snip--]
```

Figure 37.11-C: Composite numbers $n < 2^{32}$ of the form $n = 24\,j + 7$ that pass the Lucas type test. Five of them are strong pseudoprime to some base smaller than $100,000$.

Numbers of the form $n = 24\,j + 7$ satisfy the preconditions of the Lucas test *except* for the condition that $2^t > k$ where $n = k\,2^t - 1$. We test whether $H_{(n+1)/4} \equiv 0 \bmod n$, as in the Lucas test. Note that $H_n = T_n(2)$ where $T_n(x)$ is the n-th Chebyshev polynomial of the first kind. We use the fast algorithm for its computation described in section 34.2.3 on page 672 for the test routine:

```
1   bool test_7mod24(ulong n)
2   {
3       ulong nu = (n+1) >> 2;
4       umod_t t = chebyT2(nu, n);   // == chebyT(nu, 2, n);
```

```
5       return  (0==u1);
6   }
```

The function `chebyT2()` is given in [FXT: mod/chebyshev1.cc]. Figure 37.11-C gives composite numbers $n < 2^{32}$ that pass the test. The complete list of such numbers is given in [FXT: data/pseudo-7mod24.txt], there are just 64 entries. There are only five entries that are strong pseudoprimes to any basis $a < 100,000$, all shown in figure 37.11-C.

The data suggests that composites of the form $n = 24\,j + 7$ that pass the test and are pseudoprime to a small base are extremely rare. The implied test would cover 1/8 of all candidates (that are not divisible by 2 or 3), as eight numbers (1, 5, 7, 11, 13, 17, 19 and 23) are coprime to 24.

```
           n:   SPP bases a<100,000
       30739:                                              [--snip--]
      153931:                                          97917619:
      249331:                                          100079611:  4820
     1575859:                                          124134067:
     1960243:                                           [--snip--]
     2557627:  36814 49266 49267 86080                 2946282799:
     3444403:                                          3075304399:
     3767347:  26452 79860 94736                       3145717759:
     3881179:  47489 67676 72825 73841 84995 87856     3299597407:
     3882283:                                          3554502799:
    14324491:                                          3554889199:
    14970499:                                          4091977039:
    15894163:                                          4207009999:
```

Figure 37.11-D: Composite numbers $n < 2^{32}$ of the form $n = 24\,k + 19$ that pass the Lucas type test. Four of them are strong pseudoprime to some base smaller than 100,000.

For numbers of the form $n = 24\,j+19$ we use a different test: here we check whether $U_{(n+1)/4-1} \equiv 0 \bmod n$ where $U_0 = 0$, $U_1 = 1$ and $U_k = 4\,U_{k-1} - U_{k-2}$ (the Chebyshev polynomial of the second kind, $U_n(x)$, evaluated at $x = 2$). The function for testing is

```
1   bool test_19mod24(ulong n)
2   {
3       ulong nu = ((n+1) >> 2) - 1;
4       umod_t t = chebyU2(nu, n);  // == chebyU(nu, 2, n);
5       return  (0==t);
6   }
```

were the function `chebyU2()` is given in [FXT: mod/chebyshev2.cc]. The list [FXT: data/pseudo-19mod24.txt] contains all (155) composites $n < 2^{32}$ that pass the test. An extract is shown in figure 37.11-D. Just four numbers $n < 2^{32}$ are also strong pseudoprimes to any base $a < 100,000$.

The application of second order recurrent sequences to primality testing is described in [27]: define the sequence W_k by

$$W_k \;=\; P\,W_{k-1} - Q\,W_{k-2}, \qquad W_0 = 0, \quad W_1 = 1 \tag{37.11-7}$$

Then n is a *Lucas pseudoprime* (with parameters P and Q) if $W_{n\pm1} \equiv 0 \bmod n$, where the sign depends on whether $D = P^2 - 4\,Q$ is a square modulo n. For both cases considered here we have $n = 12\,j + 7$, $D = 16 - 4 = 12 = 4 \cdot 3$, and 3 is not a square modulo n. The test would be (note that $W_n = U_{n-1}(2)$)

```
1   bool lucas_7mod12(ulong n)
2   {
3       ulong nu = n;
4       umod_t t = chebyU2(nu, n);
5       return  (0==t);
6   }
```

This test is passed by far more composites than the two tests considered here. A primality test combining a Lucas-type test and a test for strong pseudoprimality has been suggested in [213]. No composite that passes the test has been found so far.

37.11.4.6 An observation about Mersenne numbers

An interesting observation about Mersenne numbers is that the following *seems* to be true:

$$M = 2^e - 1 \quad \text{prime} \quad \Longleftrightarrow \quad 3^{2^{e-1}} \equiv -3 \bmod M \tag{37.11-8}$$

Note that for odd e the condition is equivalent to $3^{(M-1)/2} \equiv -1 \bmod n$ and 3 is a non-residue. For prime exponents e it can be seen that it is very unlikely to find a composite M_e where $3^{(M-1)/2} \equiv -1 \bmod n$: the number m is a strong pseudoprime (SPP) to base 2 by construction and the right hand side of condition 37.11-8 says that m is a SPP to base 3. Given the rarity of composites that are SPP to both bases (see section 37.10.2 on page 777) the chance to find such a number among the exponentially growing Mersenne numbers is very small. Tony Reix, who observed the statement of relation 37.11-8 independently verified it for prime exponents up to 132499.

37.11.4.7 Primes that are evaluations of cyclotomic polynomials

The Mersenne numbers and Fermat numbers are actually special cases of numbers obtained as cyclotomic polynomials $Y_n(x)$ (see section 35.1.1 on page 677) evaluated at $x = 2$ (sequence A019320 of [240]). The first such numbers are shown in figure 37.11-E. The sequence of values n such that $Y_n(2)$ is prime is entry A072226 of [240]:

```
2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 13, 14, 15, 16, 17, 19, 22, 24, 26, 27, 30,
31, 32, 33, 34, 38, 40, 42, 46, 49, 56, 61, 62, 65, 69, 77, 78, 80, 85, 86,
89, 90, 93, 98, 107, 120, 122, 126, 127, 129, 133, 145, 150, 158, 165, 170,
174, 184, 192, 195, 202, 208, 234, 254, 261, ...
```

Now set $N := Y_n(2)$, testing whether $3^{(N-1)/2} \equiv -1 \bmod N$ seems to prove primality for all values of n. Note that for n a power of two the test is Pepin's test. Information about the primality of $Y_n(2)$ is given in [121]. Theorems about factorizations of $Y_n(x)$ where x is an integer are given in [129], see also [67] and [65].

The primes $Y_n(2)$ are also of interest for number theoretic transforms (see section 25.1 on page 525) because of their special structure allowing for very efficient modular reduction (see section 37.2 on page 759). A prominent example is $Y_{192}(2) = 2^{64} - 2^{32} + 1$. Note that the order of 2 modulo $Y_n(2)$ equals n.

The structure of the primes becomes (in base 10) visible if we check evaluations at 10, the first primes $Y_n(10)$ are

```
  n:   Yn(10)
  2:   11
  4:   101
 10:   9091
 12:   9901
 14:   909091
 19:   1111111111111111111
 23:   11111111111111111111111
 24:   99990001
 36:   999999000001
 38:   909090909090909091
 39:   900900900900990990991
 48:   9999999900000001
```

Finally, we do a silly thing: the factors of $Y_{2^7-1}(x)$ over GF(2) are the irreducible binary polynomials of degree 7. If we evaluate them as polynomials over $\mathbb{Z}$ at $x = 10$ and select the prime numbers we obtain 8-digit primes consisting of only zero and ones:

```
10011101    10111001    11100101    11110111    11111101
```

The same procedure, with $Y_{3^5-1}(x)$ and factoring over GF(3) gives the primes

```
101221    102101    111121    111211    112111    120011    122021
```

The list is created via

```
n=3^5-1;  f=lift(factor(polcyclo(n)*Mod(1,3)));  f=f[,1];
for(k=1, #f, v=subst(f[k],x,10); if(isprime(v), print(v)));
```

```
 n:  s=Yn(2)                                  n:  s=Yn(3)
  2:  3                                        2:  2 * 2
  3:  7                                        3:  13
  4:  5                                        4:  2 * 5
  5:  31                                       5:  11 * 11
  6:  3                                        6:  7
  7:  127                                      7:  1093
  8:  17                                       8:  2 * 41
  9:  73                                       9:  757
 10:  11                                      10:  61
 11:  23 * 89   <--= SPP [11]                 11:  23 * 3851
 12:  13                                      12:  73
 13:  8191                                    13:  797161
 14:  43                                      14:  547
 15:  151                                     15:  4561
 16:  257                                     16:  2 * 17 * 193
 17:  131071                                  17:  1871 * 34511
 18:  3 * 19                                  18:  19 * 37          <--= SPP [7]
 19:  524287                                  19:  1597 * 363889
 20:  5 * 41                                  20:  5 * 1181
 21:  7 * 337                                 21:  368089
 22:  683                                     22:  67 * 661
 23:  47 * 178481                             23:  47 * 1001523179
 24:  241                                     24:  6481
 25:  601 * 1801   <--= SPP [29]              25:  8951 * 391151
 26:  2731                                    26:  398581
 27:  262657                                  27:  109 * 433 * 8209
 28:  29 * 113                                28:  29 * 16493
 29:  233 * 1103 * 2089                       29:  59 * 28537 * 20381027
 30:  331                                     30:  31 * 271         <--= SPP [29]
 31:  2147483647                              31:  683 * 102673 * 4404047
 32:  65537                                   32:  2 * 21523361
 33:  599479                                  33:  2413941289
 34:  43691                                   34:  103 * 307 * 1021
 35:  71 * 122921                             35:  71 * 2664097031
 36:  37 * 109  <--= SPP [17, 19, 23]         36:  530713
 37:  223 * 616318177                         37:  13097927 * 17189128703
 38:  174763                                  38:  2851 * 101917
 39:  79 * 121369                             39:  13 * 313 * 6553 * 7333
 40:  61681                                   40:  42521761
 41:  13367 * 164511353                       41:  83 * 2526913 * 86950696619
 42:  5419                                    42:  7 * 43 * 2269
 43:  431 * 9719 * 2099863                    43:  431 * 380808546861411923
 44:  397 * 2113                              44:  5501 * 570461
 45:  631 * 23311      <--= SPP [5]           45:  181 * 1621 * 927001
 46:  2796203                                 46:  23535794707
 47:  2351 * 4513 * 13264529                  47:  1223 * 21997 * 5112661 * 96656723
 48:  97 * 673                                48:  97 * 577 * 769
 49:  4432676798593                           49:  491 * 4019 * 8233 * 51157 * 131713
 50:  251 * 4051                              50:  151 * 22996651
 51:  103 * 2143 * 11119                      51:  12853 * 99810171997
 52:  53 * 157 * 1613                         52:  53 * 4795973261
 53:  6361 * 69431 * 20394401                 53:  107 * 24169 * 3747607031112307667
 54:  3 * 87211                               54:  19441 * 19927
 55:  881 * 3191 * 201961                     55:  11 * 1321 * 560088668384411
 56:  15790321                                56:  430697 * 647753
 57:  32377 * 1212847                         57:  229 * 248749 * 1824179209
 58:  59 * 3033169                            58:  523 * 6091 * 5385997
 59:  179951 * 3203431780337                  59:  14425532687 * 489769993189671059
 60:  61 * 1321                               60:  47763361
 61:  2305843009213693951                     61:  603901 * 105293313660391861035901
 62:  715827883                               62:  6883 * 22434744889
 63:  92737 * 649657                          63:  1445429185300809
 64:  641 * 6700417                           64:  2 * 926510094425921
 65:  145295143558111                         65:  131 * 3701101 * 110133112994711
```

Figure 37.11-E: Evaluations s of the first cyclotomic polynomials at two (left). Entries at prime n are Mersenne numbers M_n, entries at $n = 2^k$ are Fermat numbers F_k. Composites that are strong pseudoprimes to prime bases other than two are marked with 'SPP'. The right columns show the corresponding data for evaluations at three.

37.11.4.8 Further reading

Excellent introductions into topics related to prime numbers and methods of factorization are [222], [276], and [98]. Primality tests and factorization algorithms are also described in [127]. Some of the newer factorization algorithms can be found in [90], readable surveys are [182] and [203]. Tables of factorizations of numbers of the form $b^e \pm 1$ are given in [73] which also contains much historical information.

It remains to say that a deterministic polynomial-time algorithm for proving primality was published by Agrawal, Kayal and Saxena in August 2002 [4]. While this a major breakthrough in mathematics it does not render the Rabin-Miller test worthless. Practically, 'industrial grade' primes are still produced with it, see [70]. Good introductions into the ideas behind the so-called *AKS algorithm* and its improvements are [136] and [98, p.200ff].

37.12 Complex moduli: GF(p^2)

Recall that with real numbers the equation $x^2 = -1$ has no solution, there is no real square root of minus one. The construction of complex numbers proceeds by taking pairs of real numbers $(a, b) = a + i\,b$ together with (component wise) addition $(a, b) + (c, d) = (a + c, b + d)$ and multiplication defined by $(a, b)\,(c, d) = (ac - bd, ad + bc)$. Indeed the pairs of real numbers together with addition and multiplication as given constitute a field.

We will now rephrase the construction in a way that shows how to construct an *extension field* from a given *ground field* (or *base field*). In the example above the ground field are the real numbers and the extension field are the complex numbers.

37.12.1 The construction of complex numbers

Equivalently to saying that there is no real square root of minus one we could have said that the polynomial $x^2 + 1$ has no linear factor. The construction of the complex numbers proceeds by taking numbers of the form $a + b\,i$ where i was boldly defined to be a root of the polynomial $x^2 + 1$. Now observe that if we identify $a + b\,i = b\,i + a$ with the polynomial $b\,x + a$ and use polynomial addition and multiplication modulo the polynomial $x^2 + 1$ we obtain the arithmetic of complex numbers. Addition is component wise, no modular reduction occurs. The multiplication rule can be obtained with polynomial arithmetic:

$$\begin{aligned} (b\,x + a)\;(d\,x + c) &= (bd)\,x^2 + (ad + bc)\,x + (ac) && \text{(37.12-1a)} \\ &\equiv (ad + bc)\,x + (ac - bd) \qquad (\text{mod } x^2 + 1) && \text{(37.12-1b)} \end{aligned}$$

We used the relation $x^2 = -1$, so $u\,x^2 \equiv -u \pmod{x^2 + 1}$. Identify x with i in the relations and observe that computations with a root (that is not in the ground field) of a polynomial is equivalent to computations with polynomials (whose coefficients are in the ground field) modulo the polynomial $x^2 + 1$.

When the ground field is the real numbers the story comes to an end: every polynomial of arbitrary degree n with complex coefficients has exactly n complex roots (including multiplicity). That is, we cannot use the given construction to extend the field $\mathbb{C}$ of complex numbers as every polynomial $p(x)$ with coefficients in $\mathbb{C}$ is a product of linear factors: $p(x) = (x - r_1)\,(x - r_2)\ldots,(x - r_n)$ with $r_k \in \mathbb{C}$. The field $\mathbb{C}$ is algebraically closed.

If we choose the ground field to be $\mathbb{F}_p = \text{GF}(p)$, the integers modulo a prime p, and an irreducible polynomial $c(x)$ of degree n whose coefficients are in $\mathbb{F}_p$ we get an extension field $\mathbb{F}_{p^n} = \text{GF}(p^n)$, a *finite field* with p^n elements. The special case of the *binary finite fields* GF(2^n) is treated in chapter 40 on page 875.

37.12.2 Complex finite fields

```
a= 1+1*i

a^1 =  1+1*i
a^2 =  0+2*i  // (1+x)*(1+x)   = x^2+2*x+1   == 2*x+1 - 1 = 2*x+0 == 0+2*x
a^3 =  1+2*i  // (1+x)*(0+2*x) = 2*x^2+2*x   == 2*x   - 2 = 2*x-2 == 1+2*x
a^4 =  2+0*i  // (1+x)*(1+2*x) = 2*x^2+3*x+1 == 3*x+1 - 2 = 3*x-1 == 2+0*x
a^5 =  2+2*i  //                             ^^                    ^^
a^6 =  0+1*i  //                        mod(x^2+1)            mod(3)
a^7 =  2+1*i == a^(-1)= 2+1*i
a^8 =  1+0*i == one
a^9 =  1+1*i

R=maxord==8 == Mat([2, 3])
r=ord(a)==8
R/r=1
```

Figure 37.12-A: The powers of the element $1 + 1\,x$ modulo $x^2 + 1$ and $p = 3$.

With primes of the form $p = 4k + 3$ it is possible to construct a field of complex numbers as minus one is a quadratic non-residue and so the polynomial $x^2 + 1$ is irreducible. The field is denoted by $\mathrm{GF}(p^2)$.

The rules for complex addition, subtraction and multiplication are the 'usual' ones. The field has p^2 elements of which $R = p^2 - 1$ are invertible. The maximal order equals R, so the inverse of and element a can be computed as $a^{-1} = a^{R-1} = a^{p^2-2}$.

For example, the powers of $a = 1 + x = 1 + i$ modulo $c = x^2 + 1 = 0 = 3 + 3i$ are shown in figure 37.12-A Note that the modular reduction happens with both the polynomial $x^2 + 1$ and $p = 3$. The polynomial reduction uses $x^2 = -1$.

```
a^1  = 1 + 3*x
a^2  = 2 + 2*x  // (1+3*x)*(1+3*x) = 9*x^2+6*x+1 == 6*x+1 - (9*x+9) =-3*x-8 == 2+2*x
a^3  = 1 + 2*x  // (1+3*x)*(2+2*x) = 6*x^2+8*x+2 == 8*x+2 - (6*x+6) = 2*x-4 == 1+2*x
a^4  = 0 + 4*x  // (1+3*x)*(1+2*x) = 6*x^2+5*x+1 == 5*x+1 - (6*x+6) =  -x-5 == 0+4*x
a^5  = 3 + 2*x  //                               ^^                    ^^
a^6  = 2 + 0*x  //                          mod(x^2+x+1)          mod(5)
a^7  = 2 + 1*x
a^8  = 4 + 4*x
a^9  = 2 + 4*x
a^10 = 0 + 3*x
a^11 = 1 + 4*x
a^12 = 4 + 0*x
a^13 = 4 + 2*x
a^14 = 3 + 3*x
a^15 = 4 + 3*x
a^16 = 0 + 1*x
a^17 = 2 + 3*x
a^18 = 3 + 0*x
a^19 = 3 + 4*x
a^20 = 1 + 1*x
a^21 = 3 + 1*x
a^22 = 0 + 2*x
a^23 = 4 + 1*x
a^24 = 1 + 0*x
```

Figure 37.12-B: The powers of the element $1 + 3\,x$ modulo $x^2 + x + 1$ and $p = 5$.

With primes of the form $p = 4k + 1$ it also possible to construct a field $\mathrm{GF}(p^2)$. But we have to use a different polynomial as $x^2 + 1$ is reducible modulo p and thereby the multiplication rules is different. For example, with $p = 5$ we find that $x^2 + x + 1$ is irreducible:

```
? p=5;  m=Mod(1,p)*(1+x+x^2);  polisirreducible(m)
  1
? a=Mod(1,p)*(1+3*x)
? for(k=1,p^2-1,print("a^",k," = ",component(Mod(a,m)^k,2)))
  a^1 = Mod(3, 5)*x + Mod(1, 5)
  a^2 = Mod(2, 5)*x + Mod(2, 5)
  a^3 = Mod(2, 5)*x + Mod(1, 5)
  [--snip--]
```

The complete list of powers is shown in figure 37.12-B. We see that $a = 1 + 3\,x$ has the maximal order (24), it is a primitive root. The polynomial reduction uses the relation $x^2 = -(x + 1)$.

The values of the powers of the primitive root can be used to 'randomly' fill a $p \times p$ array. With $a^k = u + x\,v$ we mark the entry at row v, column u with k:

```
[ 4 11  9 19  8]
[10  1 17 14 15]
[22  3  2  5 13]
[16 20  7 21 23]
[-- 24  6 18 12]
```

The position $0, 0$ (lower left) is not visited. Note row zero is the lowest row.

As described, the procedure fills an $p \times p$ array where p is a prime. Working with an irreducible polynomial of degree n we can fill a $p^e \times p^f \times p^g \times \ldots \times p^k$ array if $e + f + g + \ldots + k = n$: for the exponents that are one just choose one polynomial coefficient. For the exponents greater than one (say, h) combine h polynomial coefficients $c_0,\ c_1,\ \ldots,\ c_{h-1}$ and use $z_h = c_0 + c_1\,p + c_2\,p^2 + \ldots c_{h-1}\,p^{h-1}$.

37.12.3 Efficient reduction modulo certain quadratic polynomials

The polynomial $C = x^2 + 1$ is irreducible for primes of the form $4\,k + 3$ (-1 is not a square)

$$\begin{aligned}
(a\,x + b)\,(A\,x + B) &= (a\,A)\,x^2 + (a\,B + b\,A)\,x + (b\,B) && \text{(37.12-2a)}\\
&\equiv (a\,B + b\,A)\,x + (-a\,A + b\,B) \mod x^2 + 1 && \text{(37.12-2b)}\\
&= ((a + b)\,(A + B) - a\,A - b\,B)\ x + (-a\,A + b\,B) && \text{(37.12-2c)}
\end{aligned}$$

The last equality shows how to multiply two complex numbers at the cost of three real multiplications and five real additions instead of four multiplications and two additions.

The polynomial $C = x^2 + d$ is irreducible if $-d$ is not a square. We have

$$\begin{aligned}
(a\,x + b)\,(A\,x + B) &\equiv (a\,B + b\,A)\,x + (-d\,a\,A + b\,B) \mod x^2 + d && \text{(37.12-3a)}\\
&= ((a + b)\,(A + B) - a\,A - b\,B)\ x + (-d\,a\,A + b\,B) && \text{(37.12-3b)}
\end{aligned}$$

If the multiplication by d is cheap (for example, if $d = 2$) the implied technique can be a gain.

The polynomial $C := x^2 + x + 1$ has the roots $\left(-1 \pm \sqrt{-3}\right)/2$ so it is irreducible modulo p if -3 is not a square modulo p. The first few such primes p are

```
2 5 11 17 23 29 41 47 53 59 71 83 89 101 107 113 131 137 149 167 173 179 191 197 227 233 239
```

Multiplication modulo C costs only three scalar multiplications:

$$\begin{aligned}
(a\,x + b)\,(A\,x + B) &= (a\,A)\,x^2 + (a\,B + b\,A)\,x + (b\,B) && \text{(37.12-4a)}\\
&\equiv (-a\,A + a\,B + b\,A)\,x + (-a\,A + b\,B) \mod x^2 + x + 1 && \text{(37.12-4b)}\\
&= ((a - b)\,(B - A) + b\,B)\ x + (-a\,A + b\,B) && \text{(37.12-4c)}
\end{aligned}$$

For the polynomial $C = x^2 + x + d$ use

$$\begin{aligned}
(a\,x + b)\,(A\,x + B) &\equiv (-a\,A + a\,B + b\,A)\,x + (-d\,a\,A + b\,B) \mod x^2 + x + d && \text{(37.12-5a)}\\
&= ((a - b)\,(B - A) + b\,B)\ x + (-d\,a\,A + b\,B) && \text{(37.12-5b)}
\end{aligned}$$

The polynomial $C = x^2 - x - 1$ has the roots $\left(1 \pm \sqrt{5}\right)/2$, so it is irreducible modulo p if 5 is not a square modulo p. The first few such primes are:

```
2 3 7 13 17 23 37 43 47 53 67 73 83 97 103 107 113 127 137 157 163 167 173 193 197 223 227 233
```

Multiplication modulo C again costs only three scalar multiplications:

$$\begin{aligned}
(a\,x + b)\,(A\,x + B) &= (a\,A)\,x^2 + (a\,B + b\,A)\,x + (b\,B) && \text{(37.12-6a)}\\
&\equiv (a\,A + a\,B + b\,A)\,x + (a\,A + b\,B) \mod x^2 - x - 1 && \text{(37.12-6b)}\\
&= ((a + b)\,(A + B) - b\,B)\ x + (a\,A + b\,B) && \text{(37.12-6c)}
\end{aligned}$$

With the polynomial $C = x^2 - x - d$ use

$$\begin{aligned}(a\,x+b)\,(A\,x+B) &\equiv (a\,A+a\,B+b\,A)\,x+(d\,a\,A+b\,B) \mod x^2-x-d &\text{(37.12-7a)}\\ &= ((a+b)\,(A+B)-b\,B)\;x+(d\,a\,A+b\,B) &\text{(37.12-7b)}\end{aligned}$$

For polynomials of the form $C = x^2 - e\,x - d$ one has

$$\begin{aligned}(a\,x+b)\,(A\,x+B) &= (a\,A)\,x^2+(a\,B+b\,A)\,x+(b\,B) &\text{(37.12-8a)}\\ &\equiv (e\,a\,A+a\,B+b\,A)\,x+(d\,a\,A+b\,B) \mod x^2-e\,x-d &\text{(37.12-8b)}\\ &= ((a+b)\,(A+B)-[e-1]\,a\,A-b\,B)\;x+(d\,a\,A+b\,B) &\text{(37.12-8c)}\end{aligned}$$

If the multiplications by $e-1$ and d are cheap then the last equality can be useful. For example, with the polynomial $C = x^2 - 2\,x - 1$ use

$$\begin{aligned}(a\,x+b)\,(A\,x+B) &\equiv (2\,a\,A+a\,B+b\,A)\,x+(a\,A+b\,B) \mod x^2-2\,x-1 &\text{(37.12-9a)}\\ &= ((a+b)\,(A+B)-b\,B+a\,A)\;x+(a\,A+b\,B) &\text{(37.12-9b)}\end{aligned}$$

With $C = x^2 - 3\,x \pm 1$ use

$$\begin{aligned}(a\,x+b)\,(A\,x+B) &\equiv (3\,a\,A+a\,B+b\,A)\,x+(\mp a\,A+b\,B) \mod x^2-3\,x\pm 1 &\text{(37.12-10a)}\\ &= ((a+b)\,(A+B)-b\,B+2\,a\,A)\;x+(\mp a\,A+b\,B) &\text{(37.12-10b)}\end{aligned}$$

37.12.4 An algorithm for primitive 2^j-th roots

For primes with the lowest k bits set ($p \equiv (2^k - 1) \pmod{2^k}$) the largest power of two dividing the maximal order in GF(p^2) equals $N = 2^{k+1}$: $p = j\,2^k - 1$ with j odd, so $p+1 = j\,2^k$ and $p-1 = j\,2^k - 2 = 2\,(j\,2^{k-1} - 1)$, thereby $p^2 - 1 = 2^{k+1}\,[j\,(j\,2^k - 1)]$ where the term in square brackets is odd.

An algorithm for the construction of primitive 2^j-th roots in GF(p^2) for $j = 2, 3, \ldots, a$ where 2^a is the largest power of two dividing $p^2 - 1$ is given in [122] (and also in [44]):

Let $u_2 := 0$ and for $j > 2$ define

$$u_j \;:=\; \begin{cases} \Big((u_{j-1}+1)/2\Big)^{(p+1)/4} & \text{if} \quad j < a \\ \Big((u_{j-1}-1)/2\Big)^{(p+1)/4} & \text{if} \quad j = a \end{cases} \qquad \text{(37.12-11)}$$

and (for $j = 2, 3, \ldots, a$)

$$v_j \;:=\; \begin{cases} \big(+1-u_j^2\big)^{(p+1)/4} & \text{if} \quad j < a \\ \big(-1-u_j^2\big)^{(p+1)/4} & \text{if} \quad j = a \end{cases} \qquad \text{(37.12-12)}$$

where all operations are modulo p. Then $u_j + i\,v_j$ is a primitive 2^j-th root of unity in GF(p^2).

For example, with $p = 127$ (and field polynomial $x^2 + 1$) we obtain

```
j:       u_j  v_j
2:    ord(0 + i*1) = 4
3:    ord(8 + i*8) = 8
4:    ord(103 + i*21) = 16
5:    ord(68 + i*87) = 32
6:    ord(15 + i*41) = 64
7:    ord(32 + i*82) = 128
8:    ord(98 + i*38) = 256
```

For Mersenne primes $p = 2^e - 1$ one has $p^2 - 1 = (p+1)\,(p-1) = 2^e\,(2^e - 2) = 2^{e+1}\,(2^{e-1} - 1) =: 2^a\,k$ where k is odd. The highest power of two for which a primitive root exists is 2^a where $a = e + 1$ which checks with our example where $p = 127 = 2^7 - 1$.

```
p =  131071  = 2^17-1
r3 =   43811     r3^2 =  -3

k:    h + i*w =    z^(2^k)           =   h + i*r3*  u
 0:   h + i*w =    +256 + i* -56490  =   h + i*r3*    +256    ord(.) = 2^18
 1:   h + i*w =      +2 + i* +43811  =   h + i*r3*      +1    ord(.) = 2^17
 2:   h + i*w =      +7 + i* +44173  =   h + i*r3*      +4    ord(.) = 2^16
 3:   h + i*w =     +97 + i* -36933  =   h + i*r3*     +56    ord(.) = 2^15
 4:   h + i*w =  +18817 + i* +43903  =   h + i*r3*  +10864    ord(.) = 2^14
 5:   h + i*w =  -17636 + i* -35524  =   h + i*r3*  +45327    ord(.) = 2^13
 6:   h + i*w =   -5975 + i* -36232  =   h + i*r3*  +30114    ord(.) = 2^12
 7:   h + i*w =  -32446 + i* +44887  =   h + i*r3*  +58666    ord(.) = 2^11
 8:   h + i*w =  -38713 + i* -16371  =   h + i*r3*   +3123    ord(.) = 2^10
 9:   h + i*w =  +61109 + i* -46595  =   h + i*r3*  +24597    ord(.) = 2^9
10:   h + i*w =  +63110 + i* +25098  =   h + i*r3*  -48310    ord(.) = 2^8
11:   h + i*w =  +35245 + i* +14561  =   h + i*r3*   -3138    ord(.) = 2^7
12:   h + i*w =  -30756 + i* -12111  =   h + i*r3*  +50228    ord(.) = 2^6
13:   h + i*w =  -15743 + i* -35732  =   h + i*r3*  -19124    ord(.) = 2^5
14:   h + i*w =  -26425 + i* -55712  =   h + i*r3*   -1910    ord(.) = 2^4
15:   h + i*w =    -256 + i*   +256  =   h + i*r3*  +18830    ord(.) = 2^3
16:   h + i*w =       0 + i*     -1  =   h + i*r3*  +58294    ord(.) = 2^2
17:   h + i*w =      -1 + i*      0  =   h + i*r3*       0    ord(.) = 2^1
18:   h + i*w =      +1 + i*      0  =   h + i*r3*       0    ord(.) = 2^0
```

Figure 37.12-C: Elements of order 2^j in GF(p^2) where $p = 2^{17} - 1$ is a Mersenne prime.

37.12.5 Primitive 2^j-th roots with Mersenne primes

For Mersenne primes $p = 2^e - 1$ a primitive root of order 2^{e+1} (in GF(p^2) with field polynomial $x^2 + 1$) can be constructed more directly: first compute $\sqrt{-3} = 3^{(p+1)/4} = 3^{2^{e-2}}$ by squaring $e - 2$ times, then compute $1/\sqrt{2} = 2^{(e-1)/2}$ which does not require modular reduction. Now

$$z \;:=\; \frac{1}{\sqrt{2}}\left(1+\sqrt{3}\right) \;=\; \frac{1}{\sqrt{2}}\left(1+i\,\sqrt{-3}\right) \tag{37.12-13}$$

is an element of order 2^{e+1}. The number z is sometimes called the Creutzburg-Tasche primitive root as the construction was given in [99, p.200]. We have $z^2 = 2 + \sqrt{3} = 2 + i\,\sqrt{-3} = H_1 + i\,\sqrt{-3}\,U_1$ and, more general

$$z^{2^k} \;=\; H_{2^{k-1}} + i\,\sqrt{-3}\,U_{2^{k-1}} \qquad (1 \le k) \tag{37.12-14}$$

Figure 37.12-C shows the values of the successive 2^k-th powers of z in GF(p^2) where $p = 2^{17} - 1$.

The sequences $H = 2$, 7, 97, 18817, ... and $U = 1$, 4, 56, 10864, ... are those which appear in the Lucas-Lehmer test. The order of z^{2^k} is 2^{e+1-k}.

We have $H_{2j} = 2\,H_j^2 - 1$ and $H_j^2 - 3\,U_j^2 = 1$, thereby

$$H_{2j} \;=\; H_j^2 + 3\,U_j^2 \tag{37.12-15a}$$
$$U_{2j} \;=\; 2\,H_j\,U_j \tag{37.12-15b}$$

These are the index doubling formulas for the convergents of the continued fraction of $\sqrt{3}$.

37.12.6 Cosine and sine in GF(p^2)

Let z be an element of order n in GF(p^2), we would like to identify z with $\exp(2\,i\,\pi/n)$ and determine the values equivalent to $\cos(2\,\pi/n)$ and $\sin(2\,\pi/n)$. One can set

$$\cos\frac{2\,\pi}{n} \;:=\; \frac{z^2+1}{2\,z} \tag{37.12-16a}$$
$$i\,\sin\frac{2\,\pi}{n} \;:=\; \frac{z^2-1}{2\,z} \tag{37.12-16b}$$

For the choice of sin and cos the relations $\exp(x) = \cos(x) + i\,\sin(x)$ and $\sin(x)^2 + \cos(x)^2 = 1$ should hold. The first check is trivial: $\frac{z^2+1}{2\,z} + \frac{z^2-1}{2\,z} = z$. The second is also easy if we write i for some element that is the square root of -1: $(\frac{z^2+1}{2\,z})^2 + (\frac{z^2-1}{2\,z\,i})^2 = \frac{(z^2+1)^2-(z^2-1)^2}{4\,z^2} = 1$.

For the 2^j-th roots in GF(127^2) we obtain

```
0:    u+i*v=   1 + i* 0          cos=   1    i*sin=  0
1:    u+i*v= 126 + i* 0          cos= 126    i*sin=  0
2:    u+i*v=   0 + i* 1          cos=   0    i*sin=  1*i
3:    u+i*v=   8 + i* 8          cos=   8    i*sin=  8*i
4:    u+i*v= 103 + i*21          cos= 103    i*sin= 21*i
5:    u+i*v=  68 + i*87          cos=  68    i*sin= 87*i
6:    u+i*v=  15 + i*41          cos=  15    i*sin= 41*i
7:    u+i*v=  32 + i*82          cos=  32    i*sin= 82*i
8:    u+i*v=  98 + i*38          cos=  38*i  i*sin= 98
```

Note how the i swaps side with the element of highest order 2^a.

The simple construction allows us to mechanically convert fast Fourier (or Hartley) transforms with explicit trigonometric constants into the corresponding number theoretic transforms. The idea of expressing cosines and sines in terms of primitive roots was taken from [254].

37.12.7 Cosine and sine in GF(p)

What about primes of the form $p = 4k + 1$ that are used anyway for NTTs? The same construction works. The polynomial $x^2 + 1$ is reducible modulo $p = 4k + 1$, equivalently, -1 is a quadratic residue so its square root lies in GF(p). We could say: i is real modulo p if p is of the form $4k + 1$.

```
modulus= 257 == 0x101
modulus is cyclic
modulus is prime
bits(modulus)= 8.0056245  == 9 - 0.99437545
euler_phi(modulus)= 256 == 0x100 == 2^8
maxorder= 256 == 0x100
maxordelem= 3 == 0x3
max2pow= 8   (max FFT length = 2**8 == 256)
root2pow(max2pow)=3   root2pow(-max2pow)=86
sqrt(-1) =: i = 241

 8:    z=   3  = ( 173   +  87)  =  ( 173   + 107*i)
 7:    z=   9  = ( 233   +  33)  =  ( 233   +  14*i)
 6:    z=  81  = ( 123   + 215)  =  ( 123   +  99*i)
 5:    z= 136  = ( 188   + 205)  =  ( 188   + 196*i)
 4:    z= 249  = (  12   + 237)  =  (  12   + 194*i)
 3:    z=  64  = (  30   +  34)  =  (  30   +  30*i)
 2:    z= 241  = (   0   + 241)  =  (   0   +   1*i)
 1:    z= 256  = ( 256   +   0)  =  ( 256   +   0*i)
 0:    z=   1  = (   1   +   0)  =  (   1   +   0*i)
-1:    z= 256  = ( 256   +   0)  =  ( 256   +   0*i)
-2:    z=  16  = (   0   +  16)  =  (   0   + 256*i)
-3:    z= 253  = (  30   + 223)  =  (  30   + 227*i)
-4:    z=  32  = (  12   +  20)  =  (  12   +  63*i)
-5:    z= 240  = ( 188   +  52)  =  ( 188   +  61*i)
-6:    z= 165  = ( 123   +  42)  =  ( 123   + 158*i)
-7:    z= 200  = ( 233   + 224)  =  ( 233   + 243*i)
-8:    z=  86  = ( 173   + 170)  =  ( 173   + 150*i)
```

Figure 37.12-D: Roots of order 2^j modulo $p = 257 = 2^8 + 1$.

In the implementation [FXT: `class mod` in mod/mod.h] the cosine and sine values are computed from the primitive roots of order 2^j. The program [FXT: mod/modsincos-demo.cc] generates the list of 2^j-th roots and inverse roots shown in figure 37.12-D.

Again we can translate a given FFT implementation in a mechanical way.

An element modulo a prime $p = k \cdot 2^t + 1$ whose order equals 2^t can be found by the following algorithm even if the factorization of k is not known:
Choose random a where $1 < a < p-1$ and compute $s = a^k$, if $-1 = s^{2^{t-1}}$ then return s, else try another a.

The algorithm will terminate when the first element a is encountered whose order has the factor 2^t. Demonstration in pari/gp:

```
1   el2(k, t)=
2   {
3       local(p, s);
4       p=k*2^t+1;
5       for(a=2, p-2,
6           s=Mod(a,p)^k;
7           if( Mod(-1,p)==s^(2^(t-1)), return( s ); );
8       );
9   }
```

With $p = 314151729239163 \cdot 2^{26} + 1$ the algorithm terminates after testing $a = 5$ (of order $(p-1)/3$) and returning $s = 18583781386455525528042$ whose order is indeed 2^{26}.

In general, if $p = u \cdot f + 1$, $\gcd(f, u) = 1$ and $f = \prod p_i^{e_i}$ is fully factored then an element of order f can be determined by testing random values a:

1. Take a random a and set $s = a^u$.
2. If $s^{f/p_i} \neq 1$ for all prime factors p_i of f, then return s (an element of order f).
3. Go to step 1.

37.12.8 Decomposing $p = 4\,k + 1$ as sum of two squares

37.12.8.1 Direct computation

The direct way to determine u and v with $n = u^2 + v^2$ is to check (for $v = 1, 2, \ldots, \lfloor\sqrt{n}\rfloor$) whether $n - v^2$ is a perfect square. If so, return $u = \sqrt{n - v^2}$ and v:

```
1   sumofsquares_naive(k, t)=
2   { /* return [u,v] so that u^2+v^2==p =k*2^t+1 */
3       local(n, w);
4       n = k*2^t+1;
5       for (v=1, sqrtint(n),  \\ search until n-v^2 is a square
6           w = n-v^2;
7           if ( issquare(w), return( [sqrtint(w), v] ) );
8       );
9       return ( 0 );  \\ not the sum of two squares
10  }
```

The routine needs at most $\lfloor\sqrt{n}\rfloor$ steps which renders it rather useless for n large: With the prime $n = 314151729239163 \cdot 2^{26} + 1 \approx 2 \cdot 10^{22}$ we have $n = u^2 + v^2$ where $u = 132599472793$ and $v = 59158646772$ and the routine would need v steps to find the solution. The method described next finds the solution immediately.

37.12.8.2 Computation using continued fractions

The square root i of -1 can be used to find the representation of a prime $p = 4\,k + 1$ as a sum of two squares, $p = u^2 + v^2$, as follows:

1. Determine i where $i^2 = -1$ modulo p. If $i \geq p/2$ then set $i = p - i$.
2. Compute the continued fraction of p/i, it has the form $[a_0, a_1, \ldots, a_n, a_n, \ldots, a_1, a_0]$.
3. Compute the numerators of the $(n-1)$-st and the n-th convergent, P_{n-1} and P_n. Return $u = P_{n-1}$ and $v = P_n$.

Assume that $p = k \cdot 2^t + 1$ where $t \geq 2$. Use an element of order 2^t to find a square root of -1:

```
1   imag4k1(k, t)=
2   { /* determine x so that x^2=-1 modulo p=k*2^t+1 */
3       local(s);
4       s = el2(k, t);
5       s = s^(2^(t-2));
6       return( s );
7   }
```

Then

```
1   sumofsquares(k, t)=
2   { /* return [u,v] so that u^2+v^2==p =k*2^t+1 */
3       local(i, s, p, cf, q, u, v);
4       i = component( imag4k1(k, t), 2);
5       p = k*2^t+1;
6       if ( i>=p/2, i = p-i );
7       cf = contfrac(p/i);
8       cf = vector(length(cf)/2, j, cf[j]);
9       q = contfracpnqn(cf);
10      u = q[1, 1];  v = q[1, 2];
11      return( [u, v] );
12  }
```

An example, for $p = 2281$ we obtain

```
   i = 1571  \\ square root of -1
   i =  710  \\ choose smaller square root
   cf = [3, 4, 1, 2, 2, 1, 4, 3]  \\ == contfrac(2281/710)
   cf = [3, 4, 1, 2]  \\ first half on contfrac
  q = contfracpnqn(cf) =
    [45 16]  \\ == [P_4, P_3]
    [14 5]   \\ == [Q_4, Q_3] (unused)
    u=45;  v=16;  \\ u^2 + v^2 = 2025 + 256 = 2281
```

37.12.8.3 A memory saving version

An algorithm that avoids storing the continued fraction comes from the observation that u and v appear in the calculation of $\gcd(p, i)$. With

```
1   gcd_print(p, i)=
2   {
3       local( t, s );
4       if ( p<i,  t=p; p=i; i=t; );
5       s = sqrtint(p);
6       while ( i,
7           print ("  ", p, "  ", i);
8           t = p % i;  p = i;  i = t;
9       );
10  }
```

and $p = 2281$, $i = 710$ we obtain

```
     u    v
  2281  710
   710  151
   151  106
   106   45
    45   16 <--=
    16   13
    13    3
     3    1
```

The marked pair is the first where $u^2 < p$. The resulting routine is

```
1   sumofsquares_gcd(k, t)=
2   { /* return [u,v] so that u^2+v^2==p =k*2^t+1 */
3       local(s, p, i, w);
4       i = component( imag4k1(k, t), 2);
5       p = k*2^t+1;
6       if ( i>=p/2, i = p-i );
7       w = sqrtint(p);
8       while ( i,
9           if ( p<=w, return( [p,i] ) );
```

```
10          t = p % i;  p = i;  i = t;
11      );
12      return( [0,0] ); \\ failure
13  }
```

We note that by the relation $a^2+b^2=(a+i\,b)\,(a-i\,b)$ we can use the decomposition into two squares to obtain the factorization of a number over the complex integers. For example, we have $3141592653 = 3 \cdot 107 \cdot 9786893$ (over $\mathbb{Z}$) where the greatest prime factor is of the form $4k+1$. Using $9786893 = 2317^2+2102^2$ one can obtain $3141592653 = -i \cdot 3 \cdot 107 \cdot (2317+2102i) \cdot (2102+2317i)$. Pari/gp has a builtin routine for this task:

```
? factor(3141592653+0*I)
   [-I 1]   [3 1]   [107 1]   [2317 + 2102*I 1]   [2102 + 2317*I 1]
```

37.13 Solving the Pell equation

Simple continued fractions can be used to find integer solutions of the equations

$$x^2 - d\,y^2 = +1 \tag{37.13-1a}$$
$$x^2 - d\,y^2 = -1 \tag{37.13-1b}$$

Equation 37.13-1a is usually called the *Pell equation*. The name *Bhaskara equation* (or *Brahmagupta-Bhaskara equation*, used in [174]) seems more appropriate as Brahmagupta (ca. 600 AD) and Bhaskara (ca. 1100 AD) were the first to study and solve this equation.

37.13.1 Solution via continued fractions

The convergents P_k/Q_k of the continued fraction of $\sqrt{d}$ are close to $\sqrt{d}$: $(P_k/Q_k)^2 \approx d$. If we define $e_k := P_k^2 - d\,Q_k^2$ then solutions of relation 37.13-1a correspond to $e_k = +1$, solutions of 37.13-1b to $e_k = -1$.

As an example we set $d = 53$. The continued fraction of $\sqrt{53}$ is

$$\mathrm{CF}(\sqrt{53}) = [7,\ 3,1,1,3,14,\ 3,1,1,3,14,\ 3,1,1,3,14,\ \ldots] \tag{37.13-2a}$$
$$= [7,\ \overline{3,1,1,3,14}] \tag{37.13-2b}$$

We observe that the sequence is periodic after the initial term and the last term of the period is twice the initial term. Moreover, disregarding the term 14, the terms in the period form a palindrome. These properties actually hold for all simple continued fractions of square roots $\sqrt{d}$ with d not a perfect square, for the proofs see [208] or [174]. For the computation of the continued fraction of a square root a specialized version of the algorithm from section 35.3.1.2 on page 703 will be most efficient.

The table shown in figure 37.13-A gives the first convergents P_k/Q_k together with $e_k := P_k^2 - d\,Q_k^2$ ($d = 53$). The entry for $k = 4$ corresponds to the smallest solution (x, y) of $x^2-53\,y^2 = -1$: $182^2-53\cdot 25^2 = -1$. Entry $k = 9$ corresponds to $66249^2 - 53 \cdot 9100^2 = +1$, the smallest nontrivial solution to $x^2 - 53\,y^2 = +1$ (the trivial solution is $(P_{-1}, Q_{-1}) = (1, 0)$).

With $d = 19$ we obtain the continued fraction

$$\mathrm{CF}(\sqrt{19}) = [4,\ 2,1,3,1,2,8,\ 2,1,3,1,2,8,\ 2,1,3,1,2,8,\ \ldots] \tag{37.13-3a}$$
$$= [4,\ \overline{2,1,3,1,2,8}] \tag{37.13-3b}$$

with period length $l = 6$. The corresponding table (figure 37.13-B) contains solutions with $e_k = +1$ but none with $e_k = -1$:

Let e correspond the minimal nontrivial solution of $x^2 - d\,y^2 = \pm 1$, if $e = +1$ then no solution for $x^2 - d\,y^2 = -1$ exists. Nontrivial solutions with $e = +1$ always exist, solutions with $e = -1$ only exist

k :	a_k	P_k	Q_k	$e_k := P_k^2 - d\,Q_k^2$
−1:	−	0	1	+1
0:	7	7	1	−4
1:	3	22	3	+7
2:	1	29	4	−7
3:	1	51	7	+4
4:	3	182	25	−1
5:	14	2599	357	+4
6:	3	7979	1096	−7
7:	1	10578	1453	+7
8:	1	18557	2549	−4
9:	3	66249	9100	+1
10:	14	946043	129949	−4
11:	3	2904378	398947	+7
12:	1	3850421	528896	−7
13:	1	6754799	927843	+4
14:	3	24114818	3312425	−1
15:	14	344362251	47301793	+4

Figure 37.13-A: The first convergents P_k/Q_k of the continued fraction of $\sqrt{53}$.

k :	a_k	P_k	Q_k	$e_k := P_k^2 - d\,Q_k^2$
−1:	−	0	1	+1
0:	4	4	1	−3
1:	2	9	2	+5
2:	1	13	3	−2
3:	3	48	11	+5
4:	1	61	14	−3
5:	2	170	39	+1
6:	8	1421	326	−3
7:	2	3012	691	+5
8:	1	4433	1017	−2
9:	3	16311	3742	+5
10:	1	20744	4759	−3
11:	2	57799	13260	+1
12:	8	483136	110839	−3

Figure 37.13-B: The first convergents P_k/Q_k of the continued fraction of $\sqrt{19}$.

when the period length l of the continued fraction of $\sqrt{d}$ is odd. The period length l is always odd for primes of the form $p = 4\,k+1$ and never for numbers of the form $n = 4\,k+3$ and $4\,k$. If any factor f_i of d is of the form $f_i = 4\,k+3$ then no solution with $e = -1$ exists, because this would imply $x^2 \equiv -1 \bmod f_i$ but -1 is never a quadratic residue modulo $f_i = 4\,k+3$ by relation 37.8-3 on page 771.

However, all prime factors being of the form $4\,k+1$ does not guarantee that $e = -1$, the smallest examples are $205 = 5 \cdot 41$, $221 = 13 \cdot 17$, $305 = 5 \cdot 61$, and $377 = 13 \cdot 29$. The list of such numbers up to 2500 is

```
205, 221, 305, 377, 505, 545, 689, 725, 745, 793, 905, 1205, 1345, 1405, 1469, 1513,
1517, 1537, 1717, 1885, 1945, 1961, 2005, 2041, 2045, 2105, 2225, 2245, 2329, 2353
```

The sequence of numbers d with no factor of the form $4\,k+3$ such that $x^2 - d\,y^2 = -1$ has no solution is entry A031399 of [240]:

```
4, 8, 16, 20, 25, 32, 34, 40, 52, 64, 68, 80,
100, 104, 116, 128, 136, 146, 148, 160, 164, 169, 178, 194,
200, 205, 208, 212, 221, 232, 244, 256, 260, 272, 289, 292, 296, ...
```

An algorithm for compution solutions (x, y) of the equation $A\,x^2 - B\,y^2 = N$ is given in [196].

37.13.2 Powering solutions

Let $(x, y) = (P_m, Q_m)$ be a solution of relation 37.13-1b ($e_m = -1$) then $(x, y) := (P_p, Q_p)$ where

$$P_p := P_m^2 + d\,Q_m^2 \tag{37.13-4a}$$
$$Q_p := 2\,P_m\,Q_m \tag{37.13-4b}$$

is a solution of 37.13-1a ($e_p = +1$). If (P_m, Q_m) is the smallest solution with $e := e_m = -1$ then (P_p, Q_p) is the smallest solution with $e_p = e^2 = +1$. If (P_m, Q_m) is the smallest solution with $e = +1$ then (P_p, Q_p) is the second smallest solution with $e_p = +1$.

With our example from figure 37.13-A: $(P_m, Q_m) = (P_4, Q_4) = (182, 25)$ and

$$P_p := 182^2 + 53 \cdot 25^2 = 66249 = P_9 \tag{37.13-5a}$$
$$Q_p := 2 \cdot 182 \cdot 25 = 9100 = Q_9 \tag{37.13-5b}$$

If the period length of the continued fractions equals l then $m = l - 1$ and $p = 2\,l - 1$. Alternative expressions for P_p are

$$P_p = 2\,P_m^2 - e = 2\,d\,Q_m^2 + e \tag{37.13-6}$$

This can be seen as follows: $P_p + 0 = P_p + e - e = (P_m^2 + d\,Q_m) + (P_m^2 - d\,Q_m) - e$ and (second relation) $P_m^2 = d\,Q_m^2 + e$.

Further (P_t, Q_t) where

$$P_t := P_m\,\left(P_m^2 + 3\,d\,Q_m^2\right) \tag{37.13-7a}$$
$$Q_t := Q_m\,\left(3\,P_m^2 + d\,Q_m^2\right) \tag{37.13-7b}$$

is the second smallest solution with $e_t = e^3 = -1$. With our example $P_t = 24114818$ and $Q_t = 3312425$ (and $t = 14 = 3\,l - 1$). Alternative forms for P_t and Q_t are

$$P_t := P_m\,\left(4\,P_m^2 - 3\,e\right) \tag{37.13-8a}$$
$$Q_t := Q_m\,\left(4\,P_m^2 - e\right) = Q_m\,\left(4\,d\,Q_m^2 + 3\,e\right) \tag{37.13-8b}$$

Note that relations 37.13-4a and 37.13-4b are just the numerator and denominator of the second order iteration for $\sqrt{d}$ (relation 28.2-2a on page 562). Relations 37.13-7a and 37.13-7b correspond to the third order iteration (relation 28.2-2b on page 562).

If the pair (x, y) is a solution of relation 37.13-1a, then $(T_n(x), yU_{n-1}(x))$ is also a solution, where T_n and U_n are the Chebyshev polynomials of the first and second kind:

$$T_n^2(x) - d\,(y\,U_{n-1}(x))^2 = T_n^2(x) - d\,y^2\,U_{n-1}(x)^2 = \quad (37.13\text{-}9a)$$
$$T_n^2(x) - (x^2-1)\,U_{n-1}(x)^2 = 1 \quad (37.13\text{-}9b)$$

The last equality is relation 34.2-25 on page 674. Similarly, if (x, y) is a solution of 37.13-1b, then $(T_n^+(x), yU_{n-1}^+(x))$ is also a solution, by equation 34.2-30 on page 675.

37.14 Multigrades *

Let a, $a+1$ be any two successive numbers, then $a^0 = (a+1)^0$ (wow!). Let a, $a+1$, $a+2$, $a+3$ any four successive numbers, then $a^1 + (a+3)^1 = (a+1)^1 + (a+2)^1$ (still not really exciting). We write the relations as 'prototypes':

$$0^0 = 1^0 \quad (37.14\text{-}1a)$$
$$0^1 + 3^1 = 1^1 + 2^1 \quad (37.14\text{-}1b)$$

We also have

$$0^2 + 3^2 + 5^2 + 6^2 = 1^2 + 2^2 + 4^2 + 7^2 \quad (37.14\text{-}1c)$$

and

$$0^3 + 3^3 + 5^3 + 6^3 + 9^3 + 10^3 + 12^3 + 15^3 = 1^3 + 2^3 + 4^3 + 7^3 + 8^3 + 11^3 + 13^3 + 14^3 \quad (37.14\text{-}1d)$$

$$0^4 + 3^4 + 5^4 + 6^4 + 9^4 + 10^4 + 12^4 + 15^4 + 17^4 + 18^4 + 20^4 + 23^4 + 24^4 + 27^4 + 29^4 + 30^4 = \quad (37.14\text{-}1e)$$
$$= 1^4 + 2^4 + 4^4 + 7^4 + 8^4 + 11^4 + 13^4 + 14^4 + 16^4 + 19^4 + 21^4 + 22^4 + 25^4 + 26^4 + 28^4 + 31^4$$

In general, for a fixed exponent $k > 0$, let $s = 2^{k+1}$ and partition the set $\{0, 1, 2, 3, \ldots, s-1\}$ into two sets S_0, S_1 so that S_0 contains all elements j $(0 \le j < s)$ with parity zero, S_1 the elements with parity one. Then for $d \in \mathbb{C}$,

$$\sum_{a \in S_0} (a+d)^k = \sum_{b \in S_1} (b+d)^k \quad (37.14\text{-}2)$$

An example, for $k = 2$ we have $S_0 = \{0, 3, 5, 6\}$, $S_1 = \{1, 2, 4, 7\}$ and so

$$(0+d)^2 + (3+d)^2 + (5+d)^2 + (6+d)^2 = (1+d)^2 + (2+d)^2 + (4+d)^2 + (7+d)^2 \quad (37.14\text{-}3)$$

Compare to relation 37.14-1c. The prototypes for the exponent k are also valid for all positive integer exponents less than k and the exponent zero.

The equations given are special solutions of the so-called *multigrades* or *Prouhet-Tarry-Escott* problem: Find two sets A and B such that for a given k

$$\sum_{a \in A} a^j = \sum_{b \in B} b^j \qquad j = 0, 1, 2, \ldots, k \quad (37.14\text{-}4)$$

Both sets have the same number n of elements. This is a (k, n)-multigrade. We found solutions of $(k, 2^k)$-multigrades with the special property that the union of both sets contains 2^k successive numbers. The shift invariance holds again, it is a property of all multigrade solutions.

Multiplying the symbolic entries in a prototype gives a valid prototype which again is shift invariant. This means that, for example, relation 37.14-1c can be interpreted as

$$(0\,m+d)^k + (3\,m+d)^k + (5\,m+d)^k + (6\,m+d)^k = \quad (37.14\text{-}5)$$
$$= (1\,m+d)^k + (2\,m+d)^k + (4\,m+d)^k + (7\,m+d)^k$$

for $k \in \{0,1,2\}$, $d \in \mathbb{C}$ and $m \in \mathbb{C}$.

Adding and subtraction shifted versions of the prototypes gives more prototypes. We subtract from relation 37.14-1c the shifted by one version with swapped sides:

$$\begin{aligned} +\left(0^2+3^2+5^2+6^2\right) &= 1^2+2^2+4^2+7^2\right) \\ -\left(2^2+3^2+5^2+8^2\right) &= 1^2+4^2+6^2+7^2\right) \end{aligned}$$

Giving $2\cdot 2^2+8^2 = 0^2+2\cdot 6^2$. We can divide all symbolic entries by two and get:

$$2\cdot 1^2+4^2 = 0^2+2\cdot 3^2 \tag{37.14-6}$$

which really is

$$2\cdot(1\,m+d)^k+(4\,m+d)^k = (0\,m+d)^k+2\cdot(3\,m+d)^k \tag{37.14-7}$$

for $k \in \{0,1,2\}$, $d \in \mathbb{C}$ and $m \in \mathbb{C}$.

37.15 Properties of the convergents of $\sqrt{2}$ *

We give some number theoretical properties of the convergents of the continued fraction of $\sqrt{2}$.

Define $P_0 := 1$, $P_1 := 1$, $P_n := 2\,P_{n-1}+P_{n-2}$, $Q_0 := 0$, $Q_1 := 1$ and $Q_n := 2\,Q_{n-1}+Q_{n-2}$. Then P_{n+1}/Q_{n+1} is the n-th convergent of $\sqrt{2}$:

n:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
P_n:	1	3	7	17	41	99	239	577	1393	3363	8119	19601	47321	114243	275807
Q_n:	1	2	5	12	29	70	169	408	985	2378	5741	13860	33461	80782	195025

37.15.1 Pythagorean triples with difference one

```
 n :  a^2 + b^2 = c^2
 1 :  0^2 + 1^2 = 1^2
 2 :  3^2 + 4^2 = 5^2
 3 :  20^2 + 21^2 = 29^2
 4 :  119^2 + 120^2 = 169^2
 5 :  696^2 + 697^2 = 985^2
 6 :  4059^2 + 4060^2 = 5741^2
 7 :  23660^2 + 23661^2 = 33461^2
 8 :  137903^2 + 137904^2 = 195025^2
 9 :  803760^2 + 803761^2 = 1136689^2
10 :  4684659^2 + 4684660^2 = 6625109^2
11 :  27304196^2 + 27304197^2 = 38613965^2
12 :  159140519^2 + 159140520^2 = 225058681^2
13 :  927538920^2 + 927538921^2 = 1311738121^2
14 :  5406093003^2 + 5406093004^2 = 7645370045^2
15 :  31509019100^2 + 31509019101^2 = 44560482149^2
```

Figure 37.15-A: Solutions of $a^2+b^2=c^2$ with $b-a=1$.

Pythagorean triples are solutions of the Diophantine equation $a^2+b^2=c^2$. All primitive (a,b,c coprime) solutions can be obtained by choosing coprime x and y and setting $a := 2\,x\,y$, $b := x^2-y^2$, and $c := x^2+y^2$. There are Pythagorean triples where a and b differ by one, the first few are given in figure 37.15-A. The list can be obtained using $a_n = (P_{2\,n-1}-1)/2$, $b_n = a_n+1$ and $c_n = Q_{2\,n-1}$:

```
1  P(n)=( ([1,1]*[0,1;1,2]^n)[1] );
2  Q(n)=( ([0,1]*[0,1;1,2]^n)[1] );
3  pyth1(k)=
4  { /* return [a,b,c] where a^2+b^2=c^2 and b=a+1 */
5      local(p,q);
6      p = (P(2*k-1)-1)/2;
7      q = Q(2*k-1);
8      return( [p, p+1, q] );
9  }
```

Or with the equivalent

```
1   pyth2(k)=
2   { /* return [a,b,c] where a^2+b^2=c^2 and b=a+1 */
3       local(x,y);
4       x = P(k);  y=Q(k);
5       return( [x^2-y^2, 2*x*y, x^2+y^2] );
6   }
```

37.15.2 Solutions of $x^4 \pm 4 \cdot z^4 = c^2 + 1$ and $x^4 - y^4 = c^2\, d$

```
    n :  x^4 + 4*z^4 = c^2 + 1
    1 :  1^4 + 4*1^4 = 2^2 + 1
    2 :  3^4 + 4*2^4 = 12^2 + 1
    3 :  7^4 + 4*5^4 = 70^2 + 1
    4 :  17^4 + 4*12^4 = 408^2 + 1
    5 :  41^4 + 4*29^4 = 2378^2 + 1
    6 :  99^4 + 4*70^4 = 13860^2 + 1
    7 :  239^4 + 4*169^4 = 80782^2 + 1
    8 :  577^4 + 4*408^4 = 470832^2 + 1
    9 :  1393^4 + 4*985^4 = 2744210^2 + 1
   10 :  3363^4 + 4*2378^4 = 15994428^2 + 1
   11 :  8119^4 + 4*5741^4 = 93222358^2 + 1
   12 :  19601^4 + 4*13860^4 = 543339720^2 + 1
   13 :  47321^4 + 4*33461^4 = 3166815962^2 + 1
   14 :  114243^4 + 4*80782^4 = 18457556052^2 + 1
   15 :  275807^4 + 4*195025^4 = 107578520350^2 + 1
```

Figure 37.15-B: Solutions of $x^4 + 4z^4 = c^2 + 1$.

```
    n :  x^4 - 4*z^4 = c^2 + 1
    1 :  3^4 - 4*2^4 = 4^2 + 1
    2 :  17^4 - 4*12^4 = 24^2 + 1
    3 :  99^4 - 4*70^4 = 140^2 + 1
    4 :  577^4 - 4*408^4 = 816^2 + 1
    5 :  3363^4 - 4*2378^4 = 4756^2 + 1
    6 :  19601^4 - 4*13860^4 = 27720^2 + 1
    7 :  114243^4 - 4*80782^4 = 161564^2 + 1
    8 :  665857^4 - 4*470832^4 = 941664^2 + 1
    9 :  3880899^4 - 4*2744210^4 = 5488420^2 + 1
   10 :  22619537^4 - 4*15994428^4 = 31988856^2 + 1
   11 :  131836323^4 - 4*93222358^4 = 186444716^2 + 1
   12 :  768398401^4 - 4*543339720^4 = 1086679440^2 + 1
   13 :  4478554083^4 - 4*3166815962^4 = 6333631924^2 + 1
   14 :  26102926097^4 - 4*18457556052^4 = 36915112104^2 + 1
   15 :  152139002499^4 - 4*107578520350^4 = 215157040700^2 + 1
```

Figure 37.15-C: Solutions of $x^4 - 4z^4 = c^2 + 1$.

If set $a_n = (P_{2\,n} - 1)/2$, $b_n = a_n + 1$ and $c_n = Q_{2\,n}$ then we get 'almost' Pythagorean triples:

```
    n :  a^2 + b^2 = c^2 + 1
    1 :  1^2 + 2^2 = 2^2 + 1
    2 :  8^2 + 9^2 = 12^2 + 1          // b=3^2  a=2*2^2
    3 :  49^2 + 50^2 = 70^2 + 1        // a=7^2  b=2*7^2
    4 :  288^2 + 289^2 = 408^2 + 1     // b=17^2 a=2*12^2
    5 :  1681^2 + 1682^2 = 2378^2 + 1  // a=41^2 b=2*41^2
```

As the comments suggest, all equations are of the form $x^4 + 4 \cdot z^4 = c^2 + 1$. We obtain solutions via the simple routine:

```
1   deq44(k)=
2   { /* return [x,y,c] where x^4+4*z^4=c^2+1 */
3       local(x,z,c);
4       x = P(k);  z = Q(k);  c=2*x*z;
5       return( [x, z, c] );
6   }
```

The output is shown in figure 37.15-B. There are no solutions to the equation $x^4 + 4\,z^4 = c^2$ [32] so this is as close as one can get. Also, $x^4 - 4\,z^4 = c^2$ has no solutions but solutions of $x^4 - 4\,z^4 = c^2 + 1$ are obtained via

```
1   deq44m(k)=
2   { /* return [x,y,c] where x^4-4*z^4=c^2+1 */
3       local(x,z,c);
4       x = P(2*k);  z = Q(2*k);  c=2*z;
5       return( [x, z, c] );
6   }
```

The first few solutions shown in figure 37.15-C.

```
     n :  x^4 - y^4 = c^2 * d
     1 :  1^4 - 0^4 = 1^2 * 1
     2 :  4^4 - 3^4 = 5^2 * 7
     3 :  21^4 - 20^4 = 29^2 * 41
     4 :  120^4 - 119^4 = 169^2 * 239
     5 :  697^4 - 696^4 = 985^2 * 1393
     6 :  4060^4 - 4059^4 = 5741^2 * 8119
     7 :  23661^4 - 23660^4 = 33461^2 * 47321
     8 :  137904^4 - 137903^4 = 195025^2 * 275807
     9 :  803761^4 - 803760^4 = 1136689^2 * 1607521
    10 :  4684660^4 - 4684659^4 = 6625109^2 * 9369319
    11 :  27304197^4 - 27304196^4 = 38613965^2 * 54608393
    12 :  159140520^4 - 159140519^4 = 225058681^2 * 318281039
    13 :  927538921^4 - 927538920^4 = 1311738121^2 * 1855077841
    14 :  5406093004^4 - 5406093003^4 = 7645370045^2 * 10812186007
    15 :  31509019101^4 - 31509019100^4 = 44560482149^2 * 63018038201
```

Figure 37.15-D: Solutions of $x^4 - y^4 = c^2\, d$.

Finally, solutions of the equation $x^4 - y^4 = c^2 \cdot d$ (where $y = x - 1$) can be computed by

```
1   deq4421(k)=
2   { /* return [x,y,c,d] where x^4-y^4=c^2*d */
3       local(x,y,c,d);
4       y = (P(2*k-1)-1)/2;
5       x = y+1;
6       c = Q(2*k-1);  d = P(2*k-1);
7       return( [x, y, c, d] );
8   }
```

The first few are shown in figure 37.15-D. Note that equation 4 is $120^4 - 119^4 = 13^4\, 239$.

37.15.3 Triangular square numbers

A triangular number is a number of the form $n\,(n+1)/2$. There are triangular numbers that are also squares. Let z_n be the n-th triangular number that is a square. Then

$$z_n \;=\; (P_n\, Q_n)^2 \;=\; 1,\ 36,\ 1225,\ 41616,\ 1413721,\ \ldots \tag{37.15-1}$$

We can generate the z_n via

```
1   P(n)=(([1,1]*[0,1;1,2]^n)[1]);
2   Q(n)=(([0,1]*[0,1;1,2]^n)[1]);
3   zn(n)=(P(k)*Q(k))^2
```

The sequence of values z_n is entry A001110 in [240]. A recurrence for z_n is given by $z_0 = 1$, $z_1 = 36$, $z_n = 34\, z_{n-1} - z_{n-2} + 2$, a closed form expression is $z_n = \frac{1}{32}\left((3+\sqrt{8})^{2n} + (3-\sqrt{8})^{2n} - 2\right)$. A simpler recurrence gives the square roots of z_n: define $w_0 = 1$, $w_1 = 6$, and $w_n = 6\, w_{n-1} - w_{n-2}$. Then $w_n^2 = z_n$, the w_n-th square number equals z_n. The numbers w_n can be obtained as:

```
? wn(n)=return(([1,6]*[0,-1;1,6]^n)[1]);
? for(k=0,10,print1(" ",wn(k)))
 1 6 35 204 1189 6930 40391 235416 1372105 7997214 46611179
? for(k=0,10,print1(" ",wn(k)^2))
 1 36 1225 41616 1413721 48024900 1631432881 55420693056 1882672131025
```

Let $s_0 = 1$, $s_1 = 8$, $s_2 = 49$ and $s_n = 7\, s_{n-1} - 7\, s_{n-2} + s_{n-3}$. Then $s_n\,(s_n+1)/2 = z_n$, and the s_n-th triangular number equals z_n. The numbers s_n can be obtained as:

```
? sn(n)=return(([1,8,49]*[0,0,1;1,0,-7;0,1,7]^n)[1]);
? for(k=0,10,print1(" ",sn(k)))
```

```
 1 8 49 288 1681 9800 57121 332928 1940449 11309768 65918161
? tria(n)=n*(n+1)/2;
? for(k=0,10,print1(" ",tria(sn(k))))
 1 36 1225 41616 1413721 48024900 1631432881 55420693056 1882672131025
```

37.15.4 Solutions of the equations $s^3 + d^2 = w^4$

```
    n: s^3 + d^2 = w^4
    0: 1^3 + 0^2 = 1^4
    1: 8^3 + 28^2 = 6^4
    2: 49^3 + 1176^2 = 35^4
    3: 288^3 + 41328^2 = 204^4
    4: 1681^3 + 1412040^2 = 1189^4
    5: 9800^3 + 48015100^2 = 6930^4
    6: 57121^3 + 1631375760^2 = 40391^4
    7: 332928^3 + 55420360128^2 = 235416^4
    8: 1940449^3 + 1882670190576^2 = 1372105^4
    9: 11309768^3 + 63955420452028^2 = 7997214^4
   10: 65918161^3 + 2172601941851880^2 = 46611179^4
   11: 384199200^3 + 73804512448220400^2 = 271669860^4
   12: 2239277041^3 + 2507180832055219320^2 = 1583407981^4
   13: 13051463048^3 + 851703438401289936 28^2 = 9228778026^4
   14: 76069501249^3 + 2893284510097771529376^2 = 53789260175^4
   15: 443365544448^3 + 982865030016140490401 28^2 = 313506783024^4
```

Figure 37.15-E: Solutions of $s^3 + d^2 = w^4$.

The sum of cubes up to n^3 is the square of the n-th triangular number:

$$\sum_{k=1}^{n} k^3 \quad = \quad \left(\frac{n\,(n+1)}{2} \right)^2 \tag{37.15-2}$$

So the difference of the squares of two consecutive triangular numbers is a cube:

$$\left(\frac{n\,(n+1)}{2} \right)^2 - \left(\frac{(n-1)\,n}{2} \right)^2 \quad = \quad n^3 \tag{37.15-3}$$

We write $T(n)$ for the n-th triangular number. Now, we have seen that $T(s_n)$, the s_n-th triangular number, is a square: $T(s_n) = w_n^2$. Squaring gives $T(s_n)^2 = w_n^4$. Subtracting $d_n := T(s_n - 1)$ from both sides gives $s_n^3 + d_n^2 = w_n^4$. That is, the triples $[s_n,\, d_n,\, w_n]$ are a solutions of the Diophantine equation

$$s^3 + d^2 \quad = \quad w^4 \tag{37.15-4}$$

which is a somewhat surprising observation. For the computation of d_n we use $d_n = w_n^2 - s_n = z_n - s_n$:

```
1   deq(n)=
2   { /* Return [s,d,w] so that  s^3 + d^2 = w^4 */
3       local(s, t, w, d);
4       s = sn(n);
5       w = wn(n);
6       d = w^2-s;  \\ == zn(n)-s;
7       return( [s, d, w] );
8   }
```

The first few solutions of equation 37.15-4 of the given form are shown in figure 37.15-E. The equations with odd index are of the form $8\,u^6 + d^2 = w^4$. The equations with even index are of the form $u^6 + d^2 = w^4$, corresponding to Pythagorean triples giving triangles with hypothenuse a square and one side a third power. For example (equation 2): $(7^3)^2 + (1176)^2 = (35^2)^2$ which is $7^2 + 24^2 = 5^4$ multiplied by 7^4.

The triangular square numbers were determined by Euler. That triangular square numbers yield solutions to $s^3 + d^2 = w^4$ was observed 1912 by H. B. Mathieu.

We note that there exist solutions that are not of the given form. A simple search reveals some of them, we use the routine

```
1   N=200;
2   { for (s=1, N,  s3 = s^3;
3       for (w=1, N,  w4 = w^4;
4           t=w4-s3;
5           if ( issquare(t) && (t!=0),
6               d=sqrtint(t);
7               g=gcd(s,gcd(d,w));
8               print(s, "^3 + ", d, "^2 = ", w, "^4   [", g, "]");
9           );
10   ); ); }
```

All solutions where $0 < s \leq 200$ and $0 < w \leq 200$ are

```
s^3 + d^2 = w^4    [gcd(s,d,w)]
 8^3 + 28^2 = 6^4   [2]
 18^3 + 27^2 = 9^4   [9]
 23^3 + 6083^2 = 78^4   [1]
 36^3 + 63^2 = 15^4   [3]
 49^3 + 1176^2 = 35^4   [7]
 108^3 + 648^2 = 36^4   [36]
 108^3 + 3807^2 = 63^4   [9]
 126^3 + 2925^2 = 57^4   [3]
 128^3 + 1792^2 = 48^4   [16]
 135^3 + 4941^2 = 72^4   [9]
 136^3 + 4785^2 = 71^4   [1]
 143^3 + 433^2 = 42^4   [1]
```

In his amazing note [91] Cohen gives for the Diophantine equation

$$a^4 + b^3 \;=\; c^2 \tag{37.15-5}$$

the parametric solution

$$a \;=\; 6 \cdot s \cdot t \cdot (4 \cdot s^4 + 3 \cdot t^4) \tag{37.15-6a}$$
$$b \;=\; 16 \cdot s^8 - 168 \cdot s^4 \cdot t^4 + 9 \cdot t^8 \tag{37.15-6b}$$
$$c \;=\; 64 \cdot s^{12} + 1584 \cdot s^8 \cdot t^4 - 1188 \cdot s^4 \cdot t^8 - 27 \cdot t^{12} \tag{37.15-6c}$$

where $s, t \in \mathbb{Z}$.

Regarding triangular numbers we note a nice generalization of the relation

$$\left(\sum_{k=1}^{n} k \right)^2 \;=\; \sum_{k=1}^{n} k^3 \tag{37.15-7}$$

which is a consequence of

$$\sum_{k=1}^{n} k \;=\; \frac{n\,(n+1)}{2} \;=\; T(n) \tag{37.15-8a}$$
$$\sum_{k=1}^{n} k^3 \;=\; \left(\frac{n\,(n+1)}{2} \right)^2 \;=\; T(n)^2 \tag{37.15-8b}$$

Indeed, we have

$$\left(\sum_{d \backslash N} \nu(d) \right)^2 \;=\; \sum_{d \backslash N} \nu(d)^3 \tag{37.15-9}$$

where $\nu(d)$ counts the number of divisors of d. Relation 37.15-7 is the special case where $N = p^{n-1}$. The divisors of N are 1, p, p^2, ..., p^{n-1} whose numbers of divisors in turn are 1, 2, 3, ..., n.

Relation 37.15-7 can easily be proven for the case that N is a product of n distinct primes: $n = p_1 \cdot p_2 \cdot \cdot p_n$. There are $\binom{n}{j}$ divisors with exactly j prime factors, each having 2^j divisors. Thereby the sum on the left

hand side of relation 37.15-9 is equal to

$$\left(\sum_{j=0}^{n}\binom{n}{j}2^j\right)^2 \;=\; ((1+2)^n)^2 \;=\; \sum_{j=0}^{n}\binom{n}{j}2^{3j} \;=\; \sum_{j=0}^{n}\binom{n}{j}\left(2^j\right)^3 \tag{37.15-10a}$$

The quantity is always an even power of 3:

$$=\; \sum_{j=0}^{n}\binom{n}{j}8^j \;=\; (1+8)^n \;=\; 9^n \tag{37.15-10b}$$

This holds only if N factors into mutually distinct primes.

We close with two relations taken from the marvelous book of A. H. Beiler [32, p.161-162]:

$$2\left(\sum_{k=1}^{n}T(k)\right)^2 \;=\; \sum_{k=1}^{n}k^5+\sum_{k=1}^{n}k^7 \tag{37.15-11}$$

$$3\left(\sum_{k=1}^{n}T(k)\right)^3 \;=\; \sum_{k=1}^{n}T(k)^3+2\sum_{k=1}^{n}T(k)^4 \tag{37.15-12}$$

37.16 Multiplication of hypercomplex numbers *

An n-dimensional vector space (over a field) together with component-wise addition and a *multiplication table* that defines the product of any two (vector) components defines an *algebra*.

The product of two elements $x=\sum_k \alpha_k\, e_k$ and $y=\sum_j \beta_j\, e_j$ of the algebra is defined as

$$x\cdot y \;=\; \sum_{k,j=0}^{n-1}\left[(\alpha_k\cdot\beta_j)\, e_k\, e_j\right] \tag{37.16-1}$$

The quantities $e_k\, e_j$ are given in the multiplication table of the algebra. These can be arbitrary elements of the algebra, that is, linear combinations of the components. For example, a 2-dimensional algebra over the reals could have the following multiplication table:

```
          e0                  e1
e0:  (5*e1 + 3*e0)   (239*e0 + 3.1415*e1)
e1:     (0)          (17*e1 + 2.71828*e0)
```

Note that there is no neutral element of multiplication ('one'). Further, the algebra has *zero divisors*: the equation $x\cdot y=0$ has a solution where neither element is zero, namely $x=e_1$, and $y=e_0$. As almost all randomly defined algebras, it is completely uninteresting.

In what follows we will only consider algebras over the reals where the product of two components equals ± 1 times another component. For example, the *complex* numbers are a two-dimensional algebra (over the real numbers) with the multiplication table

```
        e0   e1

e0:  +e0  +e1
e1:  +e1  -e0
```

Which is, using the symbols '`1`' and '`i`',

```
       1   i

1:  +1  +i
i:  +i  -1
```

We will denote the components of an n-dimensional algebra by the numbers $0, 1, \ldots, n-1$. The multiplication table for the complex numbers would thus be written as

```
       0   1

0:  +0  +1
1:  +1  -1
```

37.16.1 The Cayley-Dickson construction

The *Cayley-Dickson construction* recursively defines multiplication tables for certain algebras where the dimension is a power of two. Let a, A, b and B be elements of a 2^{n-1}-dimensional algebra U. Define the multiplication rule for an algebra V (of dimension 2^n) written as pairs of elements of U via

$$(a,\, b) \cdot (A,\, B) \quad := \quad (a \cdot A - B \cdot b^*,\; a^* \cdot B + A \cdot b) \tag{37.16-2}$$

where the *conjugate* C^* of an element $C = (a, b)$ is defined as

$$(a,\, b)^* \;:=\; (a^*,\, -b) \tag{37.16-3}$$

and the conjugate of a real number a equals a (unmodified). The construction leads to multiplications tables where the product of two units equals plus or minus one other unit only: $e_i \cdot e_j = \pm e_k \; \forall i, j$.

```
        0   1   2   3   4   5   6   7    8   9   a   b   c   d   e   f

0:     +0  +1  +2  +3  +4  +5  +6  +7   +8  +9  +a  +b  +c  +d  +e  +f
1:     +1  -0  -3  +2  -5  +4  +7  -6   -9  +8  +b  -a  +d  -c  -f  +e
2:     +2  +3  -0  -1  -6  -7  +4  +5   -a  -b  +8  +9  +e  +f  -c  -d
3:     +3  -2  +1  -0  -7  +6  -5  +4   -b  +a  -9  +8  +f  -e  +d  -c
4:     +4  +5  +6  +7  -0  -1  -2  -3   -c  -d  -e  -f  +8  +9  +a  +b
5:     +5  -4  +7  -6  +1  -0  +3  -2   -d  +c  -f  +e  -9  +8  -b  +a
6:     +6  -7  -4  +5  +2  -3  -0  +1   -e  +f  +c  -d  -a  +b  +8  -9
7:     +7  +6  -5  -4  +3  +2  -1  -0   -f  -e  +d  +c  -b  -a  +9  +8

8:     +8  +9  +a  +b  +c  +d  +e  +f   -0  -1  -2  -3  -4  -5  -6  -7
9:     +9  -8  +b  -a  +d  -c  -f  +e   +1  -0  +3  -2  +5  -4  -7  +6
a:     +a  -b  -8  +9  +e  +f  -c  -d   +2  -3  -0  +1  +6  +7  -4  -5
b:     +b  +a  -9  -8  +f  -e  +d  -c   +3  +2  -1  -0  +7  -6  +5  -4
c:     +c  -d  -e  -f  -8  +9  +a  +b   +4  -5  -6  -7  -0  +1  +2  +3
d:     +d  +c  -f  +e  -9  -8  -b  +a   +5  +4  -7  +6  -1  -0  -3  +2
e:     +e  +f  +c  -d  -a  +b  -8  -9   +6  +7  +4  -5  -2  +3  -0  -1
f:     +f  -e  +d  +c  -b  -a  +9  -8   +7  -6  +5  +4  -3  -2  +1  -0
```

Figure 37.16-A: Multiplication table for the sedenions. The entry in row R, column C gives the product $R \cdot C$ of the components R and C (hexadecimal notation).

Figure 37.16-A gives the multiplication table for a 16-dimensional algebra, the so-called *sedenions*. The upper left (8×8) quarter gives the multiplication rule for the *octonions* (or *Cayley numbers*), the upper left 4×4 subtable gives the rule for the *quaternions* and the upper left 2×2 subtable corresponds to the complex numbers. Note that multiplication is in general neither commutative (only up to dimension 2) nor associative (only up to dimension 4).

The 2^n-dimensional algebras are (for $n > 1$) referred to as *hypercomplex numbers*. There is no generally accepted naming scheme for the algebras beyond dimension 16. We will use the names '2^n-ions'.

This form (relation 37.16-2) of the construction is given in [22], an alternative form is used in [111]:

$$(a,\, b) \cdot (A,\, B) \quad := \quad (a \cdot A - B^* \cdot b,\; b \cdot A^* + B \cdot a) \tag{37.16-4}$$

It leads to a table that is the transposed of figure 37.16-A.

```
          0 1 2 3 4 5 6 7  8 9 a b c d e f

    0:    + + + + + + + +  + + + + + + + +
    1:    + - - + - + + -  - + + - + - - +
    2:    + + - - - - + +  - - + + + + - -
    3:    + - + - - + - +  - + - + + - + -
    4:    + + + + - - - -  - - - - + + + +
    5:    + - + - + - + -  - + - + - + - +
    6:    + - - + + - - +  - + + - - + + -
    7:    + + - - + + - -  - - + + - - + +

    8:    + + + + + + + +  - - - - - - - -
    9:    + - + - + - - +  + - + - + - - +
    a:    + - - + + + - -  + - - + + + - -
    b:    + + - - + - + -  + + - - + - + -
    c:    + - - - - + + +  + - - - - + + +
    d:    + + - + - - - +  + + - + - - - +
    e:    + + + - - + - -  + + + - - + - -
    f:    + - + + - - + -  + - + + - - + -
```

Figure 37.16-B: Signs in the multiplication table for sedenions.

By construction, $e_0^2 = e_0$, $e_k^2 = -e_0$, $e_0\, e_k = e_k\, e_0 = e_k$, and $e_k\, e_j = -e_j\, e_k$ whenever both of k and j are nonzero (and $k \neq j$). Further,

$$e_k\, e_j \quad = \quad \pm e_x \quad \text{where} \quad x = k \,\text{XOR}\, j \tag{37.16-5}$$

where the sign is to be determined. Figure 37.16-B shows the pattern of the signs of the sedenion algebra. The lower left quarter is the transposed of the upper left quarter, so is the lower right quarter, except for its top row. The upper right quarter is (except for its first row) the negated upper left quarter. These observations, together with the partial antisymmetry can be cast into an algorithm to compute the signs [FXT: arith/cayley-dickson-demo.cc]:

```
1   int CD_sign_rec(ulong r, ulong c, ulong n)
2   // Signs in the multiplication table for the
3   //   algebra of n-ions (where n is a power of two)
4   //   that is obtained by the Cayley-Dickson construction:
5   // If component r is multiplied with component c then the
6   //   result is CD_sign_rec(r,c,n) * (r XOR c).
7   {
8       if ( (r==0) || (c==0) )  return +1;
9       if ( c>=r )
10      {
11          if ( c>r )   return  -CD_sign_rec(c, r, n);
12          else  return -1;  // r==c
13      }
14      // here r>c (triangle below diagonal)
15
16      ulong h = n>>1;
17      if ( c>=h )  // right
18      {
19          // (upper right unreached)
20          return   CD_sign_rec(c-h, r-h, h); // lower right
21      }
22      else  // left
23      {
24          if ( r>=h )  return   CD_sign_rec(c, r-h, h); // lower left
25          else         return   CD_sign_rec(r, c, h); // upper left
26      }
27  }
```

The function uses at most $2 \cdot \log_2(n)$ steps. Note that the second row in the table is (the signed version of) the *Thue-Morse sequence*, see section 1.16.1 on page 42. The matrix filled with entries ± 1 according to figure 37.16-B is a *Hadamard matrix*, see section 18 on page 365. The sequence of signs, read by antidiagonals, and setting $0 := +$ and $1 := -$, is entry A118685 of [240].

An iterative version of the function is [FXT: arith/cayley-dickson-demo.cc]:

```
1   inline void cp2(ulong a, ulong b, ulong &u, ulong &v)  { u=a; v=b; }
2   //
3   inline int CD_sign_it(ulong r, ulong c, ulong n)
4   {
5       int s = +1;
6
7   start:
8       if ( (r==0) || (c==0) )  return s;
9       if ( c==r )  return -s;
10      if ( c>r )   { swap2(r,c); s=-s; }
11      n >>= 1;
12      if ( c>=n )  cp2(c-n, r-n, r, c);
13      else if ( r>=n )  cp2(c, r-n, r, c);
14
15      goto start;
16  }
```

The computation of all $2^{26} = 67,108,864$ signs in the multiplication table for the '2^{13}-ions' takes less than 6 seconds (about 8 seconds with the recursive routine).

```
          0  1  2  3  4  5  6  7

    0:   +0 +1 +2 +3 +4 +5 +6 +7          + + + + + + + +
    1:   +1 -0 +6 +4 -3 +7 -2 -5          + - + + - + - -
    2:   +2 -6 -0 +7 +5 -4 +1 -3          + - - + + - + -
    3:   +3 -4 -7 -0 +1 +6 -5 +2          + - - - + + - +
    4:   +4 +3 -5 -1 -0 +2 +7 -6          + + - - - + + -
    5:   +5 -7 +4 -6 -2 -0 +3 +1          + - + - - - + +
    6:   +6 +2 -1 +5 -7 -3 -0 +4          + + - + - - - +
    7:   +7 +5 +3 -2 +6 -1 -4 -0          + + + - + - - -
```

Figure 37.16-C: Alternative multiplication table for the octonions (left) and its sign pattern (right).

An alternative multiplication table for the octonions is given in figure 37.16-C. Its sign pattern is the 8×8 Hadamard matrix shown in figure 18.1-A on page 366. Properties of this representation and the relation to shift register sequences are described in [105].

37.16.2 Fast multiplication of quaternions

```
    +--  0   1   2   3      +--  0   1   2   3        +--  0   1   2   3
    |                      |                        |
    0:   0   1   2   3      0:  -0   1   2   3      1 0:   0*  1   2   3
    1:   1   0   3   2      1:   1  -0   3   2      i 1:   1  -0   3  -2*
    2:   2   3   0   1      2:   2   3  -0   1      j 2:   2  -3* -0   1
    3:   3   2   1   0      3:   3   2   1  -0      k 3:   3   2  -1* -0
```

Figure 37.16-D: Scheme for the length-4 dyadic convolution (left), same with bucket zero negated (middle) and the multiplication table for the units of the quaternions (right). The asterisks mark those entries where the sign is different from the scheme in the middle.

```
 -0  1  2  3   4  5  6  7       0  1  2  3   4  5  6  7       #0  1  2  3   4  5  6  7
  1 -0  3  2   5  4  7  6       1 -0  3 -2   5 -4 -7  6        1  0  3 #2   5 #4 #7  6
  2  3 -0  1   6  7  4  5       2 -3 -0  1   6  7 -4 -5        2 #3  0  1   6  7 #4 #5
  3  2  1 -0   7  6  5  4       3  2 -1 -0   7 -6  5 -4        3  2 #1  0   7 #6  5 #4

  4  5  6  7  -0  1  2  3       4 -5 -6 -7  -0  1  2  3        4 #5 #6 #7   0  1  2  3
  5  4  7  6   1 -0  3  2       5  4 -7  6  -1 -0 -3  2        5  4 #7  6  #1  0 #3  2
  6  7  4  5   2  3 -0  1       6  7  4 -5  -2  3 -0 -1        6  7  4 #5  #2  3  0 #1
  7  6  5  4   3  2  1 -0       7 -6  5  4  -3 -2  1 -0        7 #6  5  4  #3 #2  1  0
```

Figure 37.16-E: Scheme for the length-8 dyadic convolution with bucket zero negated (left) and multiplication table for the octonions (middle, taken from [111]). There are 22 places where the signs differ (right, marked with '#'). This leads to an algorithm involving $8 + 22 = 30$ multiplications.

Quaternion multiplication can be achieved in eight real multiplication using the dyadic convolution (see section 22.7 on page 463): the scheme in figure 37.16-D suggests to use the dyadic convolution with bucket zero negated as a starting point which costs 4 multiplications. Some entries have to be corrected then which costs four more multiplications.

```
1       // f[] == [ re1, i1, j1, k1 ]
2       // g[] == [ re2, i2, j2, k2 ]
3       c0 := f[0] * g[0]
4       c1 := f[3] * g[2]
5       c2 := f[1] * g[3]
6       c3 := f[2] * g[1]
7
8       // length-4 dyadic convolution:
9       walsh(f[])
10      walsh(g[])
11      for i:=0 to 3  g[i] := (f[i] * g[i])
12      walsh(g[])
13
14      // normalization and correction:
15      g[0] :=   2 * c0 - g[0] / 4
16      g[1] := - 2 * c1 + g[1] / 4
17      g[2] := - 2 * c2 + g[2] / 4
18      g[3] := - 2 * c3 + g[3] / 4
```

The algorithm is taken from [150] which also gives a second variant.

The complex multiplication by three real multiplications (relation 37.12-2c on page 795) corresponds to one length-2 Walsh dyadic convolution and the correction for the product of the imaginary units:

```
1       // f[] == [ re1, im1 ]
2       // g[] == [ re2, im2 ]
3       c0 := f[1] * g[1]  // == im1 * im2
4       // length-2 dyadic convolution:
5       { f[0], f[1] } := { f[0]+f[1], f[0]-f[1] }
6       { g[0], g[1] } := { g[0]+g[1], g[0]-g[1] }
7       g[0] := f[0] * g[0]
8       g[1] := f[1] * g[1]
9       { g[0], g[1] } := { g[0]+g[1], g[0]-g[1] }
10      // normalization:
11      f[0] := f[0] / 2
12      g[0] := g[0] / 2
13      // correction:
14      g[0] := -2 * c0 + g[0]
15      // here:  g[] == [ re1 * re2 - im1 * im2,  re1 * im2 + im1 * re2 ]
```

For complex numbers of high precision multiplication is asymptotically equivalent to two real multiplications as one FFT based (complex linear) convolution can be used for the computation. Similarly, high precision quaternion multiplication is as expensive as four real multiplications. Figure 37.16-E shows an equivalent construction for the octonions leading to an algorithm with 30 multiplications.

Chapter 38

Binary polynomials

We introduce binary polynomials and the arithmetic operations on them. We further describe tests for irreducibility and primitivity. Finally, a method for the factorization of binary polynomials is given. Many of the algorithms shown can easily be implemented in hardware. The arithmetic operations with binary polynomials are the underlying methods for computations in binary finite fields which are treated in chapter 40. An important application are the linear feedback shift registers, they are described in chapter 39

A polynomial with coefficients in the field GF(2) = $\mathbb{Z}/2\mathbb{Z}$ (that is, 'coefficients modulo 2') is called a *binary polynomial*. The operations proceed as for usual polynomials except that the coefficients have to be reduced modulo two. To represent a binary polynomial in a binary computer one uses words where the bits are set at the positions where the polynomial coefficients are one. We stick to the convention that the constant term goes to the least significant bit. It turns out that the arithmetic operations can be implemented quite easily in an efficient manner.

38.1 The basic arithmetical operations

Addition of binary polynomials is the XOR operation. Subtraction is the very same operation.

Multiplication of a binary polynomial by its independent variable x is simply a shift to the left.

The following routines are given in [FXT: bpol/bitpol-arith.h].

38.1.1 Multiplication and squaring

Multiplication of two polynomials A and B is identical to the usual (binary algorithm for) multiplication, except that no carry occurs:

```
 1  inline ulong bitpol_mult(ulong a, ulong b)
 2  // Return  A * B
 3  {
 4      ulong t = 0;
 5      while ( b )
 6      {
 7          if ( b & 1 )  t ^= a;
 8          b >>= 1;
 9          a <<= 1;
10      }
11      return  t;
12  }
```

As for integer multiplication with the C-type `unsigned long`, the result will silently overflow if $\deg(A) + \deg(B)$ is equal to or greater than the word length (`BITS_PER_LONG`). If the operation `t^=a;`

```
          1..11.111 * 11.1.1
                              product as bitpol    ordinary product
           1..11.111   1    t= .....1..11.111   c= ......1..11.111
                       0
         1..11.111..   1    t= ...1.1111.1.11   c= ....11....1..11
                       0
       1..11.111....   1    t= .1.11.1..11.11   c= ..11..11.....11
      1..11.111.....   1    t= 11.....1111.11   c= 1.......11...11
```

```
          1..11.111  *  1..11.111
                              product as bitpol    ordinary product
           1..11.111   1    t= ........1..11.111   c= .........1..11.111
         1..11.111.    1    t= .......11.1.11..1   c= ........111.1..1.1
        1..11.111..    1    t= ......1111....1.1   c= ......1...1......1
                       0
       1..11.111....   1    t= ....1.1..1111.1.1   c= ....11.111111...1
      1..11.111.....   1    t= ...11..1....1.1.1   c= ...1....1.11.1...1
                       0
                       0
    1..11.111........  1    t= 1.....1.1...1.1.1   c= .1.1111..111.1...1
```

Figure 38.1-A: Multiplication (top) and squaring (bottom) of binary polynomials and numbers.

was replaced with `t+=a;` the ordinary (integer) product would be returned [FXT: gf2n/bitpolmultdemo.cc], see figure 38.1-A (top). When a binary polynomial $p = \sum_{k=0}^{d} a_k\, x^k$ is squared, the result equals $p^2 = \sum_{k=0}^{d} a_k\, x^{2k}$, figure 38.1-A (bottom). Thereby, for squaring, we just have to move the bits from position k to position $2k$:

```
1   inline ulong bitpol_square(ulong a)
2   // Return A * A
3   {
4       ulong t = 0,  m = 1UL;
5       while ( a )
6       {
7           if ( a&1 )  t ^= m;
8           m <<= 2;
9           a >>= 1;
10      }
11      return t;  // == bitpol_mult(a, a);
12  }
```

This version will unlikely give a speedup, but the equivalent function `bit_zip()` (see section 1.15 on page 39) can be a win if the degree of a is not too small.

38.1.2 Optimization of the squaring and multiplication routines

The routines for multiplication and squaring can be optimized by partially unrolling which avoids branches. As given, the function is compiled to:

```
  0: 31 c9                xor    %ecx,%ecx  // t = 0
  2: 48 85 ff             test   %rdi,%rdi  // a
  5: ba 01 00 00 00       mov    $0x1,%edx  // m = 1
  a: 74 1b                je     27 <_Z13bitpol_squarem+0x27>  // a==0 ?

 10: 48 89 c8             mov    %rcx,%rax  // tmp = t
 13: 48 31 d0             xor    %rdx,%rax  // tmp ^= m
 16: 40 f6 c7 01          test   $0x1,%dil  // if ( a&1 )
 1a: 48 0f 45 c8          cmovne %rax,%rcx  // then t = tmp
 1e: 48 c1 e2 02          shl    $0x2,%rdx  // m <<= 2
 22: 48 d1 ef             shr    %rdi       // a >>= 1
 25: 75 e9                jne    10 <_Z13bitpol_squarem+0x10>  // a!=0 ?

 27: 48 89 c8             mov    %rcx,%rax
 2a: c3                   retq
```

The `if()`-statement does not cause a branch so we unroll the contents of the loop 4-fold. Further, we move the `while()` statement to the end the loop to avoid the initial branch:

```
1   inline ulong bitpol_square(ulong a)
```

```
2   {
3       ulong t = 0,  m = 1UL;
4       do
5       {
6           if ( a&1 )  t ^= m;
7           m <<= 2;  a >>= 1;
8           if ( a&1 )  t ^= m;
9           m <<= 2;  a >>= 1;
10          if ( a&1 )  t ^= m;
11          m <<= 2;  a >>= 1;
12          if ( a&1 )  t ^= m;
13          m <<= 2;  a >>= 1;
14      }
15      while ( a );
16      return t;
17  }
```

Now we obtain much better machine code:

```
   0: 31 c9                xor    %ecx,%ecx  // t = 0
   2: ba 01 00 00 00       mov    $0x1,%edx  // m = 1

   7: 48 89 c8             mov    %rcx,%rax  // tmp = t
   a: 48 31 d0             xor    %rdx,%rax  // tmp ^= m
   d: 40 f6 c7 01          test   $0x1,%dil  // if ( a&1 )
  11: 48 0f 45 c8          cmovne %rax,%rcx  // then t = tmp
  15: 48 c1 e2 02          shl    $0x2,%rdx  // m <<= 2
  19: 48 d1 ef             shr    %rdi       // a >>= 1

  1c: 48 89 c8             mov    %rcx,%rax
  1f: 48 31 d0             xor    %rdx,%rax
  22: 40 f6 c7 01          test   $0x1,%dil
  26: 48 0f 45 c8          cmovne %rax,%rcx
  2a: 48 c1 e2 02          shl    $0x2,%rdx
  2e: 48 d1 ef             shr    %rdi

  31: 48 89 c8             mov    %rcx,%rax
  [--snip--]
  43: 48 d1 ef             shr    %rdi

  46: 48 89 c8             mov    %rcx,%rax
  [--snip--]
  58: 48 d1 ef             shr    %rdi

  5b: 75 aa                jne    7 <_Z13bitpol_squarem+0x7>  // a!=0 ?
  5d: 48 89 c8             mov    %rcx,%rax
  60: c3                   retq
```

The multiplication algorithm is optimized in the same way. For squaring one can also use the bit-zip function given in section 1.15 on page 39:

```
1   inline ulong bitpol_square(ulong a)
2   {
3       return bit_zip( a );
4   }
```

The higher half of the bits of the argument must be zero.

38.1.3 Exponentiation

With a multiplication (and squaring) function at hand, it is straightforward (see section 27.6 on page 555) to implement the algorithm for binary exponentiation:

```
1   inline ulong bitpol_power(ulong a, ulong e)
2   // Return A ** e
3   {
4       if ( 0==e )  return 1;
5
6       ulong s = a;
7       while ( 0==(e&1) )
8       {
9           s = bitpol_square(s);
10          e >>= 1;
11      }
12
13      a = s;
```

```
14      while ( 0!=(e>>=1) )
15      {
16          s = bitpol_square(s);
17          if ( e & 1 )  a = bitpol_mult(a, s);
18      }
19      return  a;
20  }
```

Note that overflow will occur even for moderate exponents.

38.1.4 Quotient and remainder

The remainder a modulo b division can be computed by initializing $A = a$ and subtracting $B = x^j \cdot b$ with $\deg(B) = \deg(A)$ from A at each step. The computation is finished as soon as $\deg b > \deg A$. As C-code:

```
1   inline ulong bitpol_rem(ulong a, ulong b)
2   // Return  R = A % B = A - (A/B)*B
3   // Must have: B!=0
4   {
5       const ulong db = highest_one_idx(b);
6       ulong da;
7       while ( db <= (da=highest_one_idx(a)) )
8       {
9           if ( 0==a )  break;  // needed because highest_one_idx(0)==highest_one_idx(1)
10          a ^= (b<<(da-db));
11      }
12      return  a;
13  }
```

The function `highest_one_idx()` is given in section 1.6 on page 16. The following version may be superior if the degree of a is small or if no fast version of the function `highest_one_idx()` is available:

```
1       while ( b <= a )
2       {
3           ulong t = b;
4           while ( (a^t) > t )  t <<= 1;
5           // =^= while ( highest_one(a) > highest_one(t) )  t <<= 1;
6           a ^= t;
7       }
8       return  a;
```

The quotient and remainder of two polynomials is computed as follows:

```
1   inline void bitpol_divrem(ulong a, ulong b, ulong &q, ulong &r)
2   // Set R, Q so that  A == Q * B + R.
3   // Must have B!=0.
4   {
5       const ulong db = highest_one_idx(b);
6       q = 0;  // quotient
7       ulong da;
8       while ( db <= (da=highest_one_idx(a)) )
9       {
10          if ( 0==a )  break;  // needed because highest_one_idx(0)==highest_one_idx(1)
11          a ^= (b<<(da-db));
12          q ^= (1UL<<(da-db));
13      }
14      r = a;
15  }
```

The division routine does the same computation but discards the remainder:

```
1   inline ulong bitpol_div(ulong a, ulong b)
2   // Return  Q = A / B
3   // Must have B!=0.
4   {
5       [--snip--]  // identical code
6       return  q;
7   }
```

38.1.5 Greatest common divisor (GCD)

The polynomial greatest common divisor (GCD) can be computed with the Euclidean algorithm [FXT: bpol/bitpol-gcd.h]:

```
inline ulong bitpol_gcd(ulong a, ulong b)
// Return  polynomial gcd(A, B)
{
    if ( a<b )  { ulong t=a; a=b; b=t; }
    // here: b<=a
    while ( 0!=b )
    {
        ulong c = bitpol_rem(a, b);
        a = b;
        b = c;
    }
    return  a;
}
```

The binary GCD algorithm can be implemented as follows:

```
inline ulong bitpol_binary_gcd(ulong a, ulong b)
{
    if ( a < b )  swap2(a, b);
    if ( b==0 )  return a;

    ulong k = 0;
    while ( !((a|b)&1) )  // both divisible by x
    {
        k++;
        a >>= 1;
        b >>= 1;
    }
    while ( !(a&1) )  a >>= 1;
    while ( !(b&1) )  b >>= 1;

    while ( a!=b )
    {
        if ( a < b )  { ulong t=a; a=b; b=t; };  // swap if deg(A)<deg(B)
        ulong t = (a^b) >> 1;
        while ( !(t&1) )  t >>= 1;
        a = t;
    }

    return  a << k;
}
```

With a fast bit-scan instruction we can optimize the function:

```
inline ulong bitpol_binary_gcd(ulong a, ulong b)
{
    { // one (or both) of a, b zero?
        ulong ta = a&b,  to = a|b;
        if ( ta==to )  return to;
    }

    ulong ka = lowest_one_idx(a);
    a >>= ka;
    ulong kb = lowest_one_idx(b);
    b >>= kb;
    ulong k = ( ka<kb ? ka : kb );

    while ( a!=b )
    {
        if ( a < b )  { ulong t=a; a=b; b=t; }  // swap if deg(A)<deg(B)
        ulong t = (a^b) >> 1;
        a = (t >> lowest_one_idx(t));
    }
    return  a << k;
}
```

Note that the comment

```
if ( a < b )  { ulong t=a; a=b; b=t; };  // swap if deg(A)<deg(B)
```

is not strictly correct as the swap can also happen with $\deg(a) = \deg(b)$ but that does no harm.

38.1.6 Exact division

Let C be a binary polynomial in x with constant term one. We use the relation (for power series)

$$\frac{1}{C} \;=\; \frac{1}{1-Y} \;=\; (1+Y)\,(1+Y^2)\,(1+Y^4)\,(1+Y^8)\,\ldots\,(1+Y^{2^n}) \mod x^{2^{n+1}} \qquad (38.1\text{-}1)$$

where $Y = 1 - C$. Now let $Y = x^{e_1} + x^{e_2} + \ldots + x^{e_k}$ where $e_i \geq 1$ and $e_{i+1} > e_i$. Then $Y^q = x^{qe_1} + x^{qe_2} + \ldots + x^{qe_k}$ whenever q is a power of two, and the multiplication by $(1 - Y^q)$ is obtained by shifts and subtractions. If A is an exact multiple of C then $R = A/C$ is a polynomial that can be computed as follows. We assume that arrays of N bits are used for the polynomials.

1. Set $R := A$ and let e_i (for $i = 1, 2, \ldots, k$) be the (ordered) positions of the nonzero coefficients of C. Set $q := 1$.

2. If $qe_1 \geq N$ then return R.

3. Set $T := 0$. For $j = 1, 2, \ldots, k$, set $T := T + R\,x^{qe_j}$. The multiplications with x^{qe_j} are left shifts by qe_j positions. Set $R := T$.

4. Set $q := 2\,q$ and goto step 2.

The method is most efficient when k, the number of nonzero coefficients of $C - 1$, is small. Sometimes one can reduce the work by dividing by $C\,D$ and finally multiplying by D for some appropriate D. For example, with all-ones polynomials $C = 1 + x + x^2 + \ldots + x^k$ and $D = 1 + x$, then $C\,D = 1 + x^{k+1}$. If C is of the form $x^u\,(1 + \ldots + x^k)$ then A/C can obviously be computed as $(A/x^u)/(C/x^u)$.

The most simple example is $C = 1 + x$ where the above procedure reduces to the inverse reversed Gray code given in section 1.16.6 on page 45:

```
inline ulong bitpol_div_xp1(ulong a)
// Return power series  A / (x+1)
// If A is a multiple of x+1, then the returned value
//   is the exact division by x+1
{
    a ^= a<<1;  // rev_gray ** 1
    a ^= a<<2;  // rev_gray ** 2
    a ^= a<<4;  // rev_gray ** 4
    a ^= a<<8;  // rev_gray ** 8
    a ^= a<<16; // rev_gray ** 16
#if  BITS_PER_LONG >= 64
    a ^= a<<32;  // for 64bit words
#endif
    return a;
}
```

For the division by $x^2 + 1$ use

```
inline ulong bitpol_div_x2p1(ulong a)
// Return power series  A / (x^2+1)
// If A is a multiple of x^2+1, then the returned value
//   is the exact division by x^2+1
{
    a ^= a<<2;  // rev_gray ** 2
    a ^= a<<4;  // rev_gray ** 4
    a ^= a<<8;  // rev_gray ** 8
    a ^= a<<16; // rev_gray ** 16
#if  BITS_PER_LONG >= 64
    a ^= a<<32;  // for 64bit words
#endif
    return a;
}
```

An algorithm for the exact division by $C = 2^k \pm 1$ (over $\mathbb{Z}$) is given in section 1.23.2 on page 60.

38.2 Multiplication for polynomials of high degree

We used the straightforward multiplication scheme whose asymptotic cost is $\sim N^2$ for polynomials of degree N. This is fine when working with polynomials of small degree. For the multiplication of two polynomials U and V both of (high, even) degree N write $U = U_0 + U_1\,x^{N/2}$, $V = V_0 + V_1\,x^{N/2}$ and use the scheme

$$U\,V \;=\; U_0 \cdot V_0\,(1 + x^{N/2}) + (U_1 - U_0) \cdot (V_0 - V_1)\,x^{N/2} + U_1 \cdot V_1\,(x^{N/2} + x^N) \tag{38.2-1}$$

recursively. Only the three multiplications indicated by a dot are expensive, the multiplications by a power of x are just shifts. The resulting scheme is the *Karatsuba multiplication for polynomials*, relation 27.2-3 on page 543 interpreted for polynomials (set $x^{N/2} = B$). Recursive application of the scheme leads to the asymptotic cost $\sim N^{\log_2(3)} \approx N^{1.585}$. When working with polynomials of high degree the implementation of the Karatsuba scheme is a must.

We give a generalization of the Karatsuba splitting that involves no constants, and several Toom-Cook schemes.

38.2.1 Splitting schemes that do not involve constants

A generalization of the Karatsuba scheme is given in [266] (see also [267]). It does not lead asymptotically better schemes than $\sim N^{\log_2(3)}$ but has a simple structure and avoids all multiplications by constants (the asymptotically better n-way splitting schemes method do involve multiplications by constants for all $n \geq 3$, see section 27.2 on page 542). We give the scheme for degree-2 (3-term) polynomials, recursive application for 3^n-term polynomials should be straightforward. Let

$$A \;=\; a_2\,x^2 + a_1\,x + a_0 \tag{38.2-2a}$$
$$B \;=\; b_2\,x^2 + b_1\,x + b_0 \tag{38.2-2b}$$
$$C \;=\; A\,B \;=\; c_4\,x^4 + c_3\,x^3 + c_2\,x^2 + c_1\,x + c_0 \tag{38.2-2c}$$

We want to compute c_0, c_1, …, c_4. With

$$d_{0,0} \;=\; a_0\,b_0 \tag{38.2-3a}$$
$$d_{1,1} \;=\; a_1\,b_1 \tag{38.2-3b}$$
$$d_{2,2} \;=\; a_2\,b_2 \tag{38.2-3c}$$
$$d_{0,1} \;=\; (a_0 + a_1)\,(b_0 + b_1) \tag{38.2-3d}$$
$$d_{0,2} \;=\; (a_0 + a_2)\,(b_0 + b_2) \tag{38.2-3e}$$
$$d_{1,2} \;=\; (a_1 + a_2)\,(b_1 + b_2) \tag{38.2-3f}$$

the c_k can be obtained as

$$c_0 \;=\; d_{0,0} \tag{38.2-4}$$
$$c_1 \;=\; d_{0,1} - d_{0,0} - d_{1,1} \tag{38.2-5}$$
$$c_2 \;=\; d_{0,2} - d_{0,0} - d_{2,2} + d_{1,1} \tag{38.2-6}$$
$$c_3 \;=\; d_{1,2} - d_{1,1} - d_{2,2} \tag{38.2-7}$$
$$c_4 \;=\; d_{2,2} \tag{38.2-8}$$

The scheme involves 6 multiplications and 13 additions. Recursive application leads to the asymptotic cost $\sim N^{\log_3(6)} \approx 1.6309$ which is slightly worse than for the 2-term scheme. However, applying this scheme first for a polynomial with $N = 3 \cdot 2^n$ terms and then using the Karatsuba scheme recursively can be advantageous.

We generalize the method for n-term polynomials and denote the scheme by KA-n. The 2-term scheme KA-2 is the Karatsuba algorithm. With

$$A = \sum_{k=0}^{n-1} a_k\, x^k \tag{38.2-9a}$$

$$B = \sum_{k=0}^{n-1} b_k\, x^k \tag{38.2-9b}$$

$$C = A\,B =: \sum_{k=0}^{2n-2} c_k\, x^k \tag{38.2-9c}$$

define

$$d_{s,s} := a_s\, b_s \qquad \text{for } s = 0,\, 1,\, \ldots,\, n-1 \tag{38.2-10a}$$

$$d_{s,t} := (a_s + b_s)\,(a_t + b_t) \quad \text{for } s + t = i,\ t > s \geq 0,\ 1 \leq i \leq 2\,n - 3 \tag{38.2-10b}$$

$$c_i^* = \sum_{\substack{s+t=i \\ t>s\geq 0}} d_{s,t} - \sum_{\substack{s+t=i \\ n-1\geq t>s\geq 0}} (d_{s,s} + d_{t,t}) \tag{38.2-10c}$$

Then

$$c_0 = d_{0,0} \tag{38.2-11a}$$

$$c_{2\,n-2} = d_{n-1,n-1} \tag{38.2-11b}$$

and for $0 < i < 2\,n - 2$:

$$c_i = \begin{cases} c_i^* & \text{if} \quad i \text{ odd} \\ c_i^* + d_{i/2,i/2} & \text{else} \end{cases} \tag{38.2-11c}$$

The Karatsuba scheme is obtained for $n = 2$.

We give pari/gp code whose output is the KA-n algorithm for given n. We need to create symbols '`ak`' (for a_k), '`bk`', and so on:

```
1   fa(k)=eval(Str("a" k))
2   fb(k)=eval(Str("b" k))
3   fc(k)=eval(Str("c" k))
4   fd(k,j)=eval(Str("d" k "" j))
```

For example, we can create a symbolic polynomial of degree 3:

```
? sum(k=0,3, fa(k) * x^k)
  a3*x^3 + a2*x^2 + a1*x + a0
```

The next routine generates the definitions of all $d_{s,t}$. It returns the number of multiplications involved:

```
1   D(n)=
2   {
3       local(mct);
4       mct = 0; \\ count multiplications
5       for (i=0, n-1,  mct+=1; print(fd(i,i), " = ", fa(i), " * ", fb(i) ) );
6       for (t=1, n-1,
7           for (s=0, t-1,
8               mct += 1;
9               print(fd(s,t), " = (", fa(s)+fa(t), ") * (", fb(s)+fb(t), ")" ) ;
10          );
11      );
12      return(mct);
13  }
```

For $n = 3$ the output is

```
d00 = a0 * b0
d11 = a1 * b1
d22 = a2 * b2
d01 = (a0 + a1) * (b0 + b1)
d02 = (a0 + a2) * (b0 + b2)
d12 = (a1 + a2) * (b1 + b2)
```

The following routine prints c_i, the coefficient of the product, in terms of several $d_{s,t}$. It returns the additions involved:

```
C(i, n)=
{
    local(N, s, act);
    act = -1; \\ count additions
    print1(fc(i), " = ");
    for (s=0, i-1,
        t = i - s;
        if ( (t>s) && (t<n),
             act += 3;
             print1(" + ", fd(s,t));
             print1(" - ", fd(s,s));
             print1(" - ", fd(t,t));
        );
    );
    if ( 0==i%2,  act+=1; print1(" + ", fd(i/2,i/2)) );
    print();
    return( act );
}
```

It has to be called for all i where $0 \le i \le 2n-2$. The algorithm is generated as follows:

```
KA(n)=
{
    local(mct, act);
    act = 0; \\ count additions
    mct = 0; \\ count multiplications
    mct = D(n);  \\ generate definitions for the d_{s,t}
    \\ generate rules for computation of c_i in terms of d_{s,t}:
    for (i=0, 2*n-2,  act+=C(i,n) );
    act += n*(n-1); \\ additions when setting up d(i,j) for i!=j
    return( [mct, act] );
}
```

With $n = 3$ we obtain

```
c0 =  + d00
c1 =  + d01 - d00 - d11
c2 =  + d02 - d00 - d22 + d11
c3 =  + d12 - d11 - d22
c4 =  + d22
```

```
d00 = a0 * b0
d11 = a1 * b1
d22 = a2 * b2
d33 = a3 * b3
d44 = a4 * b4
d01 = (a0 + a1) * (b0 + b1)
d02 = (a0 + a2) * (b0 + b2)
d12 = (a1 + a2) * (b1 + b2)
d03 = (a0 + a3) * (b0 + b3)
d13 = (a1 + a3) * (b1 + b3)
d23 = (a2 + a3) * (b2 + b3)
d04 = (a0 + a4) * (b0 + b4)
d14 = (a1 + a4) * (b1 + b4)
d24 = (a2 + a4) * (b2 + b4)
d34 = (a3 + a4) * (b3 + b4)

c0 =  + d00
c1 =  + d01 - d00 - d11
c2 =  + d02 - d00 - d22 + d11
c3 =  + d03 - d00 - d33 + d12 - d11 - d22
c4 =  + d04 - d00 - d44 + d13 - d11 - d33 + d22
c5 =  + d14 - d11 - d44 + d23 - d22 - d33
c6 =  + d24 - d22 - d44 + d33
c7 =  + d34 - d33 - d44
c8 =  + d44
```

Figure 38.2-A: Code for the algorithm KA-5.

Now we generate the definitions for the KA-5 algorithm:

```
n=5 /* n terms, degree=n-1 */
default(echo, 0);
KA(n);
```

We obtain the algorithm KA-5 shown in figure 38.2-A. The format is valid pari/gp input, so we add a few lines that print code to check the algorithm:

```
1   print("A=",sum(k=0,n-1, fa(k) * x^k))
2   print("B=",sum(k=0,n-1, fb(k) * x^k))
3   print("/* direct computation of the product: */")
4   print("C=A*B")
5   print("/* Karatsuba computation of the product: */")
6   print("K=",sum(k=0,2*n-2, fc(k) * x^k))
7   print("qq=K-C")
8   print("print( if(0==qq, \"OK.\", \" **** OUCH!\") )")
```

This gives for $n = 5$:

```
A=a4*x^4 + a3*x^3 + a2*x^2 + a1*x + a0
B=b4*x^4 + b3*x^3 + b2*x^2 + b1*x + b0
/* direct computation of the product: */
C=A*B
/* Karatsuba computation of the product: */
K=c8*x^8 + c7*x^7 + c6*x^6 + c5*x^5 + c4*x^4 + c3*x^3 + c2*x^2 + c1*x + c0
qq=K-C
print( if(0==qq, "OK.", " **** OUCH!") )
```

We can feed the output into another pari/gp session to verify the algorithm. We use the option '`-f`' that prevents colorization of the output which would confuse the verification process, the option '`-q`' suppresses the output of the version:

```
gp -f -q < karatsuba-n.gp  |  gp
```

We obtain (shortened and comments added):

```
/* definitions of d(s,t): */
b0*a0
b1*a1
b2*a2
b3*a3
b4*a4
(b0 + b1)*a0 + (a1*b0 + b1*a1)
(b0 + b2)*a0 + (a2*b0 + b2*a2)
[--snip--]
(b3 + b4)*a3 + (a4*b3 + b4*a4)

/* the c_i in terms of d(s,t), evaluated: */
b0*a0
b1*a0 + a1*b0
b2*a0 + (a2*b0 + b1*a1)
b3*a0 + (a3*b0 + (b2*a1 + a2*b1))
b4*a0 + (a4*b0 + (b3*a1 + (a3*b1 + b2*a2)))
b4*a1 + (a4*b1 + (b3*a2 + a3*b2))
b4*a2 + (a4*b2 + b3*a3)
b4*a3 + a4*b3
b4*a4

/* polynomials: */
a4*x^4 + a3*x^3 + a2*x^2 + a1*x + a0
b4*x^4 + b3*x^3 + b2*x^2 + b1*x + b0

/* direct computation of product: */
b4*a4*x^8 + (b4*a3 + a4*b3)*x^7 + (b4*a2 + (a4*b2 + b3*a3))*x^6 + [...]

/* Karatsuba computation of product: */
b4*a4*x^8 + (b4*a3 + a4*b3)*x^7 + (b4*a2 + (a4*b2 + b3*a3))*x^6 + [...]

/* difference: */
0

OK.  /* we are fine! */
```

The number of multiplications with algorithm KA-n is $(n^2 + n)/2$ which is suboptimal except for $n = 2$. However, recursive application can be worthwhile. One should start with the biggest prime factors as the number of additions is then minimized. The number of multiplications does not depend on the order of recursion (see [266] which also tabulates the number of additions and multiplications for $n \leq 128$).

With n just below a highly composite number one should add (zero-valued) 'dummy' terms and recursively use KA-n algorithms for small n. For example, with polynomials of degree 63 recursion with KA-2 (and $n = 64$) will beat the scheme "KA-7, then KA-3".

One could write code generators that create expanded versions of the recursions for n the product of small primes. When the cost of multiplication is much higher than for addition (as for binary polynomial

multiplication on general purpose CPUs) substantial savings can be expected.

38.2.2 Toom-Cook algorithms for binary polynomials

We give a 3-way and a 4-way splitting scheme taken from [49].

38.2.3 3-way splitting

```
A = a2*Y^2 + a1*Y + a0
B = b2*Y^2 + b1*Y + b0

S3 = a2 + a1 + a0;
S2 = b2 + b1 + b0;
S1 = S3 * S2;          \\ Mult (1)
S0 = a2*x^2 + a1*x;
S4 = b2*x^2 + b1*x;
S3 += S0;
S2 += S4;
S0 += a0;
S4 += b0;
S3 *= S2;              \\ Mult (2)
S2 = S0 * S4;          \\ Mult (3)
S4 = a2 * b2;          \\ Mult (4)
S0 = a0 * b0;          \\ Mult (5)

S3 += S2;

S2 += S0;  S2 /= x;  S2 += S3;
T = S4;  T *= (x^3+1);  \\ temporary variable
S2 += T;  S2 /= (x+1);  \\ exact division

S1 += S0;
S3 += S1;  S3 /= x;  S3 /= (x+1);  \\ exact division
S1 += S4;  S1 += S2;
S2 += S3;

P = S4*Y^4 + S3*Y^3 + S2*Y^2 + S1*Y + S0;
Mod(1,2)* (P - A*B) \\ == zero
```

Figure 38.2-B: Implementation of the 3-way multiplication scheme for binary polynomials. The five expensive multiplications are commented with '`Mult (n)`'.

For the multiplication of two polynomials A and B both of degree $3N$ write

$$A \;=\; a_0 + a_1\,x^N + a_2\,x^{2N} \;=:\; a_0 + a_1\,Y + a_2\,Y^2 \tag{38.2-12}$$

and identically for B. A 3-way splitting scheme for multiplication is shown in figure 38.2-B. The multiplications and divisions by x are shifts and the exact divisions are linear operations if we use the method of section 38.1.6 on page 820.

38.2.4 4-way splitting

For the multiplication of two polynomials A and B both of degree $4N$ write

$$A \;=\; a_0 + a_1\,Y + a_2\,Y^2 + a_3\,Y^3 \tag{38.2-13}$$

where $Y := x^N$, and identically for B. The 4-way splitting multiplication scheme is shown in figure 38.2-C.

```
A = a3*Y^3 + a2*Y^2 + a1*Y + a0;
B = b3*Y^3 + b2*Y^2 + b1*Y + b0;

S1 = a3 + a2 + a1 + a0;
S2 = b3 + b2 + b1 + b0;
S3 = S1 * S2;                \\ Mult (1)
S0 = a1 + x*(a2 + x*a3);
S6 = b1 + x*(b2 + x*b3);
S4 = (S0 + a3*(x+1))*x + S1;
S5 = (S6 + b3*(x+1))*x + S2;
S0 = S0*x + a0;
S6 = S6*x + b0;
S5 = S5 * S4;                \\ Mult (2)
S4 = S0 * S6;                \\ Mult (3)
S0 = a0*x^3 + a1*x^2 + a2*x;
S6 = b0*x^3 + b1*x^2 + b2*x;
S1 = S1 + S0 + a0*(x^2+x);
S2 = S2 + S6 + b0*(x^2+x);
S0 = S0 + a3;
S6 = S6 + b3;
S1 = S1 * S2;                \\ Mult (4)
S2 = S0 * S6;                \\ Mult (5)
S6 = a3 * b3;                \\ Mult (6)
S0 = a0 * b0;                \\ Mult (7)

S1 = S1 + S2 + S0*(x^4+x^2+1);
S5 = (S5 + S4 + S6*(x^4+x^2+1) + S1) \ (x^4+x);
S2 = S2 + S6 + S0*x^6;
S4 = S4 + S2 + S6*x^6 + S0;
S4 = (S4 + S5*(x^5+x)) \ (x^4+x^2);
S3 = S3 + S0 + S6;
S1 = S1 + S3;
S2 = S2 + S1*x + S3*x^2;
S3 = S3 + S4 + S5;
S1 = (S1 + S3*(x^2+x)) \ (x^4+x);
S5 = S5 + S1;
S2 = (S2 + S5*(x^2+x)) \ (x^4+x^2);
S4 = S4 + S2;

P = S6*Y^6 + S5*Y^5 + S4*Y^4 + S3*Y^3 + S2*Y^2 + S1*Y + S0;
Mod(1,2)*(P - A*B)  \\ == zero
```

Figure 38.2-C: Implementation of the 4-way multiplication scheme for binary polynomials. The seven expensive multiplications are commented with '`Mult (n)`'.

38.2.5 Multiplication methods for binary polynomials of high degrees

For polynomials of very high degree FFT-based algorithms can be used. The most simple method is to use integer multiplication without the carry phase (which is polynomial multiplication!). Our multiplication example can be recomputed using decimal digits. The carry phase of the integer multiplication is replaced by a reduction modulo 2:

```
 100110111 * 110101
== 11022223331211  // integer multiplication
== 11000001111011  // parity of digits
```

The scheme will work for polynomials of degree less than nine only. When using an FFT multiplication scheme (see section 27.3 on page 550) we can multiply polynomials up to degree N as long as the integer values 0, 1, 2 ... $N+1$ can be distinguished after computing the FFT. This is hardly a limitation at all: with the C-type `float` (24 bit mantissa) polynomials up to degree one million can be multiplied assuming at least 20 bits are correct after the FFT. With type `double` (53-bit mantissa) there is no practical limit. While the algorithm is very easy to implement it is not competitive to well implemented splitting schemes and the FFT method described in [233] or the multiplication algorithm given in [76]. An excellent source for multiplication algorithms for binary polynomials is [71].

38.3 Modular arithmetic with binary polynomials

Here we consider arithmetic of binary polynomials modulo a binary polynomial. Addition and subtraction are again the XOR operation. The functions shown in this section are given in [FXT: bpol/bitpolmod-arith.h].

38.3.1 Multiplication and squaring

Multiplication of a polynomial A by x modulo (a polynomial) C is achieved by shifting left and subtracting C if the coefficient shifted out is one:

```
1   static inline ulong bitpolmod_times_x(ulong a, ulong c, ulong h)
2   // Return  (A * x) mod C
3   // where A and C represent polynomials over Z/2Z:
4   //  W = pol(w) =: \sum_k{ [bit_k(w)] * x^k}
5   //
6   // h needs to be a mask with one bit set:
7   //  h == highest_one(c) >> 1  == 1UL << (degree(C)-1)
8   {
9       ulong s = a & h;
10      a <<= 1;
11      if ( s )  a ^= c;
12      return  a;
13  }
```

In order to avoid the repeated computation of the highest set bit we introduced the auxiliary variable h that has to be initialized as described in the comment. Section 1.6 on page 16 gives algorithms for the function `highest_one()`. Note that h needs to be computed only if the degree of the modulus C changes, which is usually only once for a series of calculations. By using the variable h we can obtain the correct result even if the degree of C equals the number of bits in a word in which case C does not fit into a word.

Multiplication of two polynomials a and b modulo C can be achieved by adding a reduction step to the binary multiplication routine:

```
1   inline ulong bitpolmod_mult(ulong a, ulong b, ulong c, ulong h)
2   // Return  (A * B) mod C
3   {
4       ulong t = 0;
5       while ( b )
6       {
7           if ( b & 1 )  t ^= a;
8           b >>= 1;
9
10          ulong s = a & h;
11          a <<= 1;
12          if ( s )  a ^= c;
13      }
14      return  t;
15  }
```

38.3.2 Optimization of the squaring and multiplication routines

Squaring a can be achieved as the multiplication $a \cdot a$. If many squaring have to be done with a fixed modulus then a optimization using a precomputed table of the residues $x^{2k} \bmod C$ shown in section 40.1 on page 875 can be useful. Squaring of the polynomial $\sum_{k=0}^{d} a_k\, x^k$ can be achieved via the computation the sum $\sum_{k=0}^{d} a_k\, x^{2k}$ modulo C. We use the auxiliary function

```
1   static inline ulong bitpolmod_times_x2(ulong a, ulong c, ulong h)
2   // Return  (A * x * x) mod C
3   {
4       { ulong s=a&h; a<<=1; if (s) a^=c; }
5       { ulong s=a&h; a<<=1; if (s) a^=c; }
6       return  a;
7   }
```

The squaring function, with a the 4-fold unrolled loop, is

```
1   static inline ulong bitpolmod_square(ulong a, ulong c, ulong h)
2   // Return A*A mod C
3   {
4       ulong t = 0, s = 1;
5       do
6       {
7           if (a&1) t^=s; a>>=1; s=bitpolmod_times_x2(s, c, h);
8           if (a&1) t^=s; a>>=1; s=bitpolmod_times_x2(s, c, h);
9           if (a&1) t^=s; a>>=1; s=bitpolmod_times_x2(s, c, h);
10          if (a&1) t^=s; a>>=1; s=bitpolmod_times_x2(s, c, h);
11      }
12      while ( a );
13      return t;
14  }
```

Whether the unrolled code is used can be specified via the line

```
#define  MULT_UNROLL  // define to unroll loops 4-fold
```

The optimization used for the multiplication routine is also unrolling as described in section 38.1.2 on page 816:

```
1   static inline ulong bitpolmod_mult(ulong a, ulong b, ulong c, ulong h)
2   {
3       ulong t = 0;
4       do
5       {
6           { if(b&1) t^=a; b>>=1; ulong s=a&h; a<<=1; if(s) a^=c; }
7           { if(b&1) t^=a; b>>=1; ulong s=a&h; a<<=1; if(s) a^=c; }
8           { if(b&1) t^=a; b>>=1; ulong s=a&h; a<<=1; if(s) a^=c; }
9           { if(b&1) t^=a; b>>=1; ulong s=a&h; a<<=1; if(s) a^=c; }
10      }
11      while ( b );
12      return  t;
13  }
```

It turns out that squaring via multiplication is slightly faster than via the described sum computation.

38.3.3 Exponentiation

A routine for modular exponentiation can be obtained using the right-to-left powering algorithm from section 27.6.1 on page 555:

```
1   inline ulong bitpolmod_power(ulong a, ulong e, ulong c, ulong h)
2   // Return (A ** e)  mod C
3   {
4       if ( 0==e )  return 1;  // avoid hang with e==0 in next while()
5
6       ulong s = a;
7       while ( 0==(e&1) )
8       {
9           s = bitpolmod_square(s, c, h);
10          e >>= 1;
11      }
12
13      a = s;
14      while ( 0!=(e>>=1) )
15      {
16          s = bitpolmod_square(s, c, h);
17          if ( e & 1 )  a = bitpolmod_mult(a, s, c, h);
18      }
19      return  a;
20  }
```

The left-to-right powering algorithm given in section 27.6.2 on page 556 can be implemented as:

```
1   inline ulong bitpolmod_power(ulong a, ulong e, ulong c, ulong h)
2   {
3       ulong s = a;
4       ulong b = highest_one(e);
5       while ( b>1 )
6       {
7           b >>= 1;
```

```
8           s = bitpolmod_square(s, c, h); // s *= s;
9           if ( e & b )  s = bitpolmod_mult(s, a, c, h);   // s *= a;
10      }
11      return s;
12  }
```

Computing a power of x can be optimized with this scheme:

```
1   inline ulong bitpolmod_xpower(ulong e, ulong c, ulong h)
2   // Return (x ** e)  mod C
3   {
4       ulong s = 2;  // 'x'
5       ulong b = highest_one(e);
6       while ( b>1 )
7       {
8           b >>= 1;
9           s = bitpolmod_square(s, c, h); // s *= s;
10          if ( e & b )  s = bitpolmod_times_x(s, c, h);   // s *= x;
11      }
12      return s;
13  }
```

38.3.4 Division by x

Division by x is possible if the modulus has a nonzero constant term (that is, $\gcd(C, x) = 1$):

```
1   static inline ulong bitpolmod_div_x(ulong a, ulong c, ulong h)
2   // Return  (A / x) mod C
3   // C must have nonzero constant term: (c&1)==1
4   {
5       ulong s = a & 1;
6       a >>= 1;
7       if ( s )
8       {
9           a ^= (c>>1);
10          a |= h;  // so it also works for  n == BITS_PER_LONG
11      }
12      return  a;
13  }
```

If we do not insist on correct results for the case that the degree of C equals the number of bits in a word, we could simply use the following two-liner:

```
if ( a & 1 )  a ^= c;
a >> 1;
```

The operation needs only about two CPU cycles. The inverse of x can be computed with:

```
1   static inline ulong bitpolmod_inv_x(ulong c, ulong h)
2   // Return  (1 / x) mod C
3   // C must have nonzero constant term: (c&1)==1
4   {
5       ulong a = (c>>1);
6       a |= h;  // so it also works for  n == BITS_PER_LONG
7       return  a;
8   }
```

38.3.5 Extended GCD, computation of the inverse, and division

The algorithm for the computation of extended GCD (EGCD) is taken from [168] [FXT: bpol/bitpol-gcd.h]:

```
1   inline ulong bitpol_egcd(ulong u, ulong v, ulong &iu, ulong &iv)
2   // Return u3 and set u1,v1 so that  gcd(u,v) == u3 == u*u1 + v*u2
3   {
4       ulong u1 = 1,  u2 = 0;
5       ulong v1 = 0,  v3 = v;
6       ulong u3 = u,  v2 = 1;
7       while ( v3!=0 )
8       {
9           ulong q = bitpol_div(u3, v3);  // == u3 / v3;
```

```
10
11          ulong t1 = u1 ^ bitpol_mult(v1, q);  // == u1 - v1 * q;
12          u1 = v1;  v1 = t1;
13
14          ulong t3 = u3 ^ bitpol_mult(v3, q);  // == u3 - v3 * q;
15          u3 = v3;  v3 = t3;
16
17          ulong t2 = u2 ^ bitpol_mult(v2, q);  // == u2 - v2 * q;
18          u2 = v2;  v2 = t2;
19      }
20
21      iu = u1;  iv = u2;
22      return u3;
23  }
```

The routine can be optimized using `bitpol_divrem()`: remove the lines

```
        ulong q = bitpol_div(u3, v3);  // == u3 / v3;
    [--snip--]
        ulong t3 = u3 ^ bitpol_mult(v3, q);  // == u3 - v3 * q;
```

and insert at the beginning of the body of the loop:

```
        ulong q, t3;
        bitpol_divrem(u3, v3, q, t3);
```

The routine computes the GCD g and two additional quantities i_u and i_v so that

$$g \quad = \quad u \cdot i_u + v \cdot i_v \tag{38.3-1}$$

When $g = 1$ we have

$$1 \quad \equiv \quad u \cdot i_u \bmod v \tag{38.3-2}$$

That is, i_u is the inverse of u modulo v. Thereby [FXT: bpol/bitpolmod-arith.h]:

```
1   inline ulong bitpolmod_inverse(ulong a, ulong c)
2   // Returns the inverse of A modulo C if it exists, else zero.
3   // Must have deg(A) < deg(C)
4   {
5       ulong i, t;  // t unused
6       ulong g = bitpol_egcd(a, c, i, t);
7       if ( g!=1 )  i = 0;
8       return i;
9   }
```

Modular division is obtained by multiplication with the inverse:

```
1   inline ulong bitpolmod_divide(ulong a, ulong b, ulong c, ulong h)
2   // Return a/b modulo c.
3   // Must have: gcd(b,c)==1
4   {
5       ulong i = bitpolmod_inverse(b, c);
6       a = bitpolmod_mult(a, i, c, h);
7       return a;
8   }
```

When working modulo prime m the inverse of a number a can be obtained as $a^{-1} = a^{m-2}$ ($m-1$ is the maximal order of an element in $\mathbb{Z}/m\mathbb{Z}$). With an irreducible (see section 38.4 on the facing page) polynomial C of degree n the inverse modulo C of a polynomial A can be obtained by computing $A^{-1} = A^{2^n-2}$ ($2^n - 1$ is the maximal order modulo C, see section 38.5 on page 834):

```
1   inline ulong bitpolmod_inverse_irred(ulong a, ulong c, ulong h)
2   // Return (A ** -1)  mod C
3   // Must have: C irreducible.
4   {
5       ulong r1 = (h<<1) - 2;  // max order minus one
6       ulong i =  bitpolmod_power(a, r1, c, h);
7       return  i;
8   }
```

38.4 Irreducible polynomials

A polynomial is called *irreducible* if it has no non-trivial factors (trivial factors are the constant polynomial '1' and the polynomial itself). A polynomial that has a non-trivial factorization is called *reducible*. A polynomial with zero constant term is always reducible because it has the factor x, except if the polynomial equals x. A binary polynomial that is irreducible has at least one term of odd degree and an odd number of terms in total.

The irreducible polynomials are the 'primes' among the polynomials.

Whether a given polynomial can be factorized does depend on its coefficient field: $p(x) = x^2 + 1$ does not factor as an 'ordinary' polynomial (coefficients in $\mathbb{R}$ or $\mathbb{Z}$) while the factorization over $\mathbb{C}$ is $(x^2+1) = (x+i)\,(x-i)$. As a binary polynomial, the factorization is $(x^2+1) = (x+1)^2$.

38.4.1 Testing for irreducibility

Irreducibility tests for binary polynomials use the fact that the polynomial $x^{2^k} + x$ has all irreducible polynomials of degree k as a factor. For example,

$$\begin{aligned} x^{2^6} + x &= x^{64} + x \\ &= (x) \cdot (x+1) \cdot \\ &\quad \cdot \left(x^2 + x + 1\right) \cdot \\ &\quad \cdot \left(x^3 + x + 1\right) \cdot \left(x^3 + x^2 + 1\right) \cdot \\ &\quad \cdot \left(x^6 + x + 1\right) \cdot \left(x^6 + x^3 + 1\right) \cdot \left(x^6 + x^4 + x^2 + x + 1\right) \cdot \\ &\quad \cdot \left(x^6 + x^4 + x^3 + x + 1\right) \cdot \left(x^6 + x^5 + 1\right) \cdot \left(x^6 + x^5 + x^2 + x + 1\right) \cdot \\ &\quad \cdot \left(x^6 + x^5 + x^3 + x^2 + 1\right) \cdot \left(x^6 + x^5 + x^4 + x + 1\right) \cdot \left(x^6 + x^5 + x^4 + x^2 + 1\right) \end{aligned} \tag{38.4-1}$$

A bit more can be said about the polynomial $x^{2^k} + x$ used in the tests: if

$$z = x^{2^d} + x \tag{38.4-2a}$$

$$s = 1 + \sum_{k=0}^{d-1} x^{2^k} \tag{38.4-2b}$$

$$t = 0 + \sum_{k=0}^{d-1} x^{2^k} = s - 1 = z/s \tag{38.4-2c}$$

then s has all degree-d irreducible polynomials as factors where the coefficient of x^{d-1} is one and the factorization of t consists of polynomials where that coefficient is zero. As an example consider the case $d = 7$. We use the notation $[a, b, c, \ldots, h] := x^a + x^b + x^c + \ldots + x^h$:

```
z = [128,1]  = s * t

s = [64,32,16,8,4,2,1,0] =          t = [64,32,16,8,4,2,1] =
    [1,0] *                             [1] *
    [7,6,0] *                           [7,1,0] *
    [7,6,3,1,0] *                       [7,3,0] *
    [7,6,4,1,0] *                       [7,3,2,1,0] *
    [7,6,4,2,0] *                       [7,4,0] *
    [7,6,5,2,0] *                       [7,4,3,2,0] *
    [7,6,5,3,2,1,0] *                   [7,5,2,1,0] *
    [7,6,5,4,0] *                       [7,5,3,1,0] *
    [7,6,5,4,2,1,0] *                   [7,5,4,3,0] *
    [7,6,5,4,3,2,0]                     [7,5,4,3,2,1,0]
```

38.4.1.1 The Ben-Or test for irreducibility

A polynomial C of degree d is reducible if $\gcd(x^{2^k} + x \bmod C, C) \neq 1$ for any $k < d$. We compute $u_k = x^{2^k}$ (modulo C) for each $k < d$ by successive squarings and test whether $\gcd(u_k + x, C) = 1$ for all k. But as a factor of degree f implies another one of degree $d - f$ it suffices to do the first $\lfloor d/2 \rfloor$ of the tests. The implied algorithm is called the *Ben-Or irreducibility test*. A C++ implementation is given in [FXT: bpol/bitpol-irred-ben-or.cc]:

```
1   bool bitpol_irreducible_q(ulong c, ulong h)
2   // Return whether C is irreducible (via the Ben-Or irreducibility test_;
3   // h needs to be a mask with one bit set:
4   //  h == highest_one(C) >> 1  == 1UL << (degree(C)-1)
5   {
6       if ( c<4 )
7       {
8           if ( c>=2 )  return true;   // x, and 1+x are irreducible
9           else         return false;  // constant polynomials are reducible
10      }
11
12      if ( 0==(1&c) )  return false;  // x is a factor
13
14      // if ( 0==(c & 0xaaaaaaaaUL ) )  return 0; // at least one odd degree term
15      // if ( 0==parity(c) )  return 0; // need odd number of nonzero coeff.
16      // if ( 0!=bitpol_test_squarefree(c) )  return 0; // must be square free
17
18      ulong d = h >> 1;
19      ulong u = 2;  // =^= x
20      while ( 0 != d )  // floor( degree/2 ) times
21      {
22          // Square r-times for coefficients of c in GF(2^r).
23          // We have r==1:
24          u = bitpolmod_square(u, c, h);
25
26          ulong upx = u ^ 2;  // =^= u+x
27
28          ulong g = bitpol_binary_gcd(upx, c);
29
30          if ( 1!=g )  return  false;   // reducible
31
32          d >>= 2;
33      }
34      return  true;  // irreducible
35  }
```

Commented out at the beginning are a few tests to check for trivial necessary conditions for irreducibility. For the test `bitpol_test_squarefree` (for a square factor) see section 38.12.2 on page 854. The routine will fail if $\deg c =$`BITS_PER_LONG`, because the gcd-computation fails in this case.

38.4.1.2 Rabin's test for irreducibility

Rabin's algorithm to test for irreducibility (given in [217, p.7]) can be stated as follows: A binary polynomial C of degree d is irreducible only if

$$x^{2^d} \equiv x \bmod C \tag{38.4-3a}$$

and, for all prime divisors p_i of d where $p_i < d$,

$$\gcd\left(x^{2^{d/p_i}} - x \bmod C, C\right) = 1 \tag{38.4-3b}$$

The number of GCD computations equals the number of prime divisors of d minus one, for d prime no GCD needs to be computed.

When the prime divisors are processed in decreasing order the successive exponents are increasing and the power of x can be updated via squarings (the idea can be found in [123]). Thereby the total number of squarings equals d which is minimal.

A C++ implementation of Rabin's test is given in [FXT: bpol/bitpol-irred-rabin.cc]. A table of auxiliary bitmasks gives the number of squarings between the GCD computations:

```
static const ulong rabin_tab[] =
{
 0UL,   // x = 0  (bits: ..........)  OPS:
 0UL,   // x = 1  (bits: ..........)  OPS:  finally sqr 1 times
 0UL,   // x = 2  (bits: ..........)  OPS:  finally sqr 2 times
 0UL,   // x = 3  (bits: ..........)  OPS:  finally sqr 3 times
 4UL,   // x = 4  (bits: ........1..)  OPS: sqr 2 times,    finally sqr 2 times
 0UL,   // x = 5  (bits: ..........)  OPS:  finally sqr 5 times
 12UL,  // x = 6  (bits: .......11..)  OPS: sqr 2 times,   sqr 1 times,   finally sqr 3 times
 0UL,   // x = 7  (bits: ..........)  OPS:  finally sqr 7 times
 16UL,  // x = 8  (bits: ......1....)  OPS: sqr 4 times,    finally sqr 4 times
 8UL,   // x = 9  (bits: .......1...)  OPS: sqr 3 times,    finally sqr 6 times
 36UL,  // x = 10 (bits: .....1..1..)  OPS: sqr 2 times,   sqr 3 times,   finally sqr 5 times
  [--snip--]
```

The testing routine is

```
inline bool bitpol_irreducible_rabin_q(ulong c, ulong h)
// Return whether C is irreducible (via Rabin's irreducibility test).
// h needs to be a mask with one bit set:
//  h == highest_one(C) >> 1  == 1UL << (degree(C)-1)
{
    if ( c<4 )
    {
        if ( c>=2 )  return true; // x, and 1+x are irreducible
        else         return false; // constant polynomials are reducible
    }

    if ( 0==(1&c) )  return false; // x is a factor

    ulong d = 1 + lowest_one_idx(h);  // degree
    ulong rt = rabin_tab[d];
    ulong m = 2; // =^= 'x'

    while ( rt > 1 )
    {
        do
        {
            --d;
            m = bitpolmod_square(m, c, h);
            rt >>= 1;
        }
        while ( 0 == (rt & 1) );

        ulong g = bitpol_binary_gcd( m ^ 2UL, c );
        if ( g!=1 )  return false;
    }

    do  { m = bitpolmod_square(m, c, h); }  while ( --d );
    if ( m ^ 2UL )  return false;

    return true;
}
```

Rabin's test will be faster than the Ben-Or test if the polynomial is irreducible. When the polynomial is reducible and has small factors (as often the case with 'random' polynomials) then the Ben-Or test will terminate faster. A comparison of the tests is given in [123].

38.4.1.3 Testing for irreducibility using less GCD computations

Call a binary polynomial C of degree n that has no linear factors and for which

$$x^{2^n} \equiv x \bmod C \tag{38.4-4}$$

holds, a *pseudo-irreducible* (PI). To be a PI a polynomial must have a factorization of the form

$$C = \left(x^{d_1} + \ldots\right)^{e_1} \left(x^{d_2} + \ldots\right)^{e_2} \ldots \left(x^{d_f} + \ldots\right)^{e_f} \tag{38.4-5}$$

where every d_j divides n and $\mathrm{lcm}\,(d_1,\, d_2, \ldots,\, d_f) = n$. Note that an irreducible polynomial is a PI.

Now suppose that C is a square of an irreducible polynomial, $C = S^2$. Then $x^{2^k} \bmod C$ has (for $k > 0$) only coefficients at even positions and can therefore not equal x. So a pure square can never be a PI. Also if C contains any square factor then it cannot be a PI. That is, $e_j = 1$ for all j.

For C to be a PI there must be an integer partition $n = \sum_j d_j$ into divisors where $l := \mathrm{lcm}\,(d_1,\, d_2, \ldots,\, d_f) = n$. Now if $l < n$ then

$$x^{2^l} \quad \equiv \quad x \bmod C \tag{38.4-6}$$

which can easily be tested. This motivates our next definition. We call a polynomial C a *strong pseudo-irreducible* (SPI) if it is a PI (relation 38.4-4) and

$$x^{2^k} \quad \not\equiv \quad x \bmod C \qquad \text{for} \quad 0 < k < n \tag{38.4-7}$$

As sets we have $\{I\} \subseteq \{SPI\} \subseteq \{PI\}$ where we wrote $\{I\}$ for the set of irreducible polynomials. We know that $\{I\} = \{PI\}$ for prime degrees. For degrees which are a prime power we have $\{I\} = \{SPI\}$, because $l < n$ unless C is irreducible.

A first version of the improved test for C of degree n and n having prime factors p_j can be given as

1. If C has a linear factor (x or $x + 1$) then return false.
2. For $k = 1, \ldots, n$ compute $s_k := x^{2^k} \bmod C$ by successive squarings. If $s_k = x$ and $k < n$ then return false. If $k = n/p_j$ for some prime divisor of n then store s_k.
3. If $s_n \neq x$ then return false.
4. For each stored s_k compute $g_k := \gcd(s_k, C)$, if $g_k \neq 1$ then return false.
5. Return true.

Now we list the degrees $1 < n < 63$ such that $\{I\} \neq \{SPI\}$ (the degrees where one needs GCD computations in general):

```
   12, 18, 20, 24, 28, 30, 36, 40, 42,     45, 48, 50,     54, 56, 60, 63  // I!=SPI
1, 12, 18, 20, 24, 28, 30, 36, 40, 42, 44, 45, 48, 50, 52, 54, 56, 60, 63  // A102467
```

The sequence is a subsequence of entry A102467 of [240], numbers such that there is no partition $n = \sum m_j\, p_j$ into prime divisors p_j with $m_j > 0$. If no such partition exists then no partition into divisors of n can exist. So if n is not in sequence A102467 we can avoid all GCD computations. We would unnecessarily compute GCDs for $n = 44$ and $n = 52$. Again, we disallowed linear factors: without this condition the degrees 6, 21, 35, 44, 52, and 55 would appear in our list.

Whether a number n appears in A102467 can be decided easily. Let $\prod_{1 \le j \le f} p_j^{e_j}$ be the factorization of n. If $f = 1$ (n is a prime or a prime power) then n is not in the list, if $f \geq 3$ (more than three prime factors) then n is in the list, and if $f = 2$ then n is in the list if and only if $\sum_j e_j \geq 3$. For $2 < n < 10^7$ there are 7,430,987 values that need GCDs (n appears in A102467), and 2,569,012 that don't.

Our second improvement of the test is to firstly check whether n appears in A102467 and if so, use the test above, else $\{I\} = \{SPI\}$ for degree n so just test whether C is a SPI.

38.5 Primitive polynomials

Let C be an irreducible polynomial. Then the sequence $p_k = x^k \bmod (C)$, $k = 1, 2, \ldots$ is periodic and the (smallest) period m of the sequence is the order of x modulo C. We call m the *period* (or *order*) of the polynomial C. For a binary polynomial of degree n the maximal period equals $2^n - 1$. If the period is maximal then all nonzero polynomials of degree $n - 1$ occur in the sequence p.

For the period m of C we have $x^m = 1 \bmod C$, so $x^m - 1 = 0 \bmod C$. That is, C divides $x^m - 1$ but no polynomial $x^k - 1$ with $k < m$.

A polynomial is called *primitive* if its period is maximal. Then the powers of x generate all nonzero binary polynomials of degree $n - 1$. The polynomial x is a generator ('primitive root') modulo C. Primitivity implies irreducibility, the converse is not true.

The situation is somewhat parallel to the operations modulo an integer:

- Among those integers m that are prime some have the primitive root 2: the sequence 2^k for $k = 1, 2, \ldots, m - 1$ contains all nonzero numbers modulo m (see chapter 25 on page 525).
- Among those polynomials C that are irreducible some are primitive: the sequence x^k for $k = 1, 2, \ldots, 2^n - 1$ contains all nonzero polynomials modulo C.

Note that there is another notion of the term 'primitive', that of a polynomial for which the greatest common divisor of all coefficients is one.

38.5.1 Roots of primitive polynomials have maximal order

A different characterization of primitivity is as follows. Suppose you want to do computations with linear combinations $A = \sum_{k=0}^{n-1} a_k \, \alpha^k$ (where $a_k \in \mathrm{GF}(2)$) of the powers of an (unknown!) root α of an irreducible polynomial $C = x^n + \sum_{k=0}^{n-1} c_k \, x^k$.

When multiplying A with the root α we obtain a term α^n which we want to get rid of. But we have

$$\alpha^n \quad = \quad \sum_{k=0}^{n-1} c_k \, \alpha^k \tag{38.5-1}$$

as α is a root of the polynomial C. Therefore we can use exactly the same modular reduction as with polynomial computation modulo C. The same is true for the multiplication of two linear combinations (of the powers of the same root α).

We see that the order of a polynomial p is the order of its root α modulo p, and that a polynomial is primitive if and only if all of its roots have maximal order.

38.5.2 Testing for primitivity

Checking a degree-d binary polynomial for primitivity by directly using the definition costs proportional 2^n operations which is prohibitive except for tiny d.

Improvements come from the multiplicative structure of the problem: it suffices to check whether $x^k \equiv 1 \bmod C$ for all divisors of the maximal order that are greater than n. However, this soon gets impractical quickly as n grows. Already with $n = 144$ one has maximal order $m = 2^{144} - 1$ (which has 262144 divisors of which 262112 have to be tested) so the computation gets expensive. Note that the implied algorithm needs the factorization of $m = 2^n - 1$.

A much better solution is a modification of the algorithm to determine the oder in a finite field given on page 766. The implementation given here uses the pari/gp language:

```
1  polorder(p) =
2  /* Order of x modulo p (p irreducible over GF(2)) */
3  {
4    local(g, g1, te, tp, tf, tx);
5    g = 'x;
6    p *= Mod(1,2);
7    te = nn_;
8    for(i=1, np_,
9      tf = vf_[i];  tp = vp_[i];  tx = vx_[i];
```

```
10      te = te / tf;
11      g1 = Mod(g, p)^te;
12      while ( 1!=g1,
13        g1 = g1^tp;
14        te = te * tp;
15      );
16    );
17    return( te );
18  }
```

The function uses the following global variables that must be set up before call:

```
1   nn_ = 0;  /* max order  =  2^n-1 */
2   np_ = 0;  /* number of primes in factorization */
3   vp_ = []; /* vector of primes */
4   vf_ = []; /* vector of factors (prime powers) */
5   vx_ = []; /* vector of exponents */
```

As given, the algorithm will do n_p exponentiations modulo p where n_p is the number of different primes in the factorization in m. For $n = 144$ one has $n_p = 17$. An implementation of the algorithm in C++ is given in [FXT: bpol/bitpol-order.cc].

A shortcut that makes the algorithm terminate as soon as the computed order drops below maximum is

```
1   polmaxorder_q(p) =
2   /* Whether order of x modulo p is maximal  (p irreducible over GF(2)) */
3   /* Early-out variant */
4   {
5     local(g, g1, te, tp, tf, tx, ct);
6     g = 'x;
7     p *= Mod(1,2);
8     te = nn_;
9     for(i=1, np_,
10      tf = vf_[i];  tp = vp_[i];  tx = vx_[i];
11      te = te / tf;
12      g1 = Mod(g, p)^te;
13      ct = 0;
14      while ( 1!=g1,
15        g1 = g1^tp;
16        te = te * tp;
17        ct = ct + 1;
18      );
19      if ( ct<tx,  return(0) );
20    );
21    return(1);
22  }
```

Range	time	# irred. tests	tests/sec
90 .. 99:	<1 sec	258	>300
190 .. 199:	6 sec	588	100
290 .. 299:	10 sec	474	58
390 .. 399:	30 sec	810	27

Determination of irreducible polynomials.

Range	time	# irred. tests	# prim. tests
90 .. 99:	4 sec	313	13
190 .. 199:	40 sec	1371	25
290 .. 299:	90 sec	864	14
390 .. 399:	260 sec	1529	14

Determination of primitive polynomials.

Figure 38.5-A: Timings and number of tests involved in the determination of 10 low-bit irreducible (top) and primitive (bottom) binary polynomials in selected ranges of their degrees. The small number of tests for primitivity indicates that an irreducible polynomial is likely primitive.

Using `polmaxorder_q()` and pari's built-in `polisirreducible()` the search for the lowest-bit primitive polynomials up to degree $n = 100$ is a matter of about 10 seconds. Extending the list up to $n = 200$ takes

3 minutes. The computation of all polynomials up to degree $n = 400$ takes less than on hour. Figure 38.5-A shows approximate timings with the determination of irreducible and primitive polynomials.

Again, the algorithm depends an precomputed factorizations. The table [FXT: data/mersenne-factors.txt] taken from [73] was used in order to save computation time.

For prime $m = 2^n - 1$ (that is, m is a Mersenne prime) irreducibility suffices for primality: The one-liner

```
x=127; for(z=1,x-1,if(polisirreducible(Mod(1,2)+t^z+t^x),print1(" ",z)))
```

finds all primitive trinomials for exponents of Mersenne primes in no time:

```
  89:  38 51
 127:  1 7 15 30 63 64 97 112 120 126
 521:  32 48 158 168 353 363 473 489
 607:  105 147 273 334 460 502
```

The computation (for $d = 607$) takes about a minute. Note we did not exploit the symmetry (reversed polynomials are also primitive). Techniques to find primitive trinomials whose degree are very big Mersenne exponents are described in [69].

Here is a surprising theorem: Let $p(x) = \sum_{k=0}^{d} c_k\, x^k$ be an irreducible binary polynomial, and $L_p(x) := \sum_{k=0}^{d} c_k\, x^{2^k}$. Then all irreducible factors of $L_p(x)/x$ (a polynomial of degree $2^d - 1$) are of degree equal to $\mathrm{ord}(p)$ (the order of x modulo $p(x)$). Especially, if $p(x)$ is primitive, then $L_p(x)/x$ is irreducible. The theorem is proven in [281] (also in [185, p.110]). An example: $x^7 + x + 1$ is primitive, so $x^{127} + x + 1$ is irreducible. But, as $2^{127} - 1$ is prime, $x^{127} + x + 1$ is also primitive. Thereby $x^{2^{127}-1} + x + 1$ is irreducible.

38.6 The number of irreducible and primitive polynomials

n:	I_n	n:	I_n	n:	I_n	n:	I_n
1:	2	11:	186	21:	99858	31:	69273666
2:	1	12:	335	22:	190557	32:	134215680
3:	2	13:	630	23:	364722	33:	260300986
4:	3	14:	1161	24:	698870	34:	505286415
5:	6	15:	2182	25:	1342176	35:	981706806
6:	9	16:	4080	26:	2580795	36:	1908866960
7:	18	17:	7710	27:	4971008	37:	3714566310
8:	30	18:	14532	28:	9586395	38:	7233615333
9:	56	19:	27594	29:	18512790	39:	14096302710
10:	99	20:	52377	30:	35790267	40:	27487764474

Figure 38.6-A: The number of irreducible binary polynomials for degrees $n \le 40$.

n:	P_n	n:	P_n	n:	P_n	n:	P_n
1:	1	11:	176	21:	84672	31:	69273666
2:	1	12:	144	22:	120032	32:	67108864
3:	2	13:	630	23:	356960	33:	211016256
4:	2	14:	756	24:	276480	34:	336849900
5:	6	15:	1800	25:	1296000	35:	929275200
6:	6	16:	2048	26:	1719900	36:	725594112
7:	18	17:	7710	27:	4202496	37:	3697909056
8:	16	18:	7776	28:	4741632	38:	4822382628
9:	48	19:	27594	29:	18407808	39:	11928047040
10:	60	20:	24000	30:	17820000	40:	11842560000

Figure 38.6-B: The number of primitive binary polynomials for degrees $n \le 40$.

The number of irreducible binary polynomials of degree n is

$$I_n \;=\; \frac{1}{n}\sum_{d\backslash n}\mu(d)\,2^{n/d} \;=\; \frac{1}{n}\sum_{d\backslash n}\mu(n/d)\,2^{d} \tag{38.6-1}$$

The Möbius function μ is defined by relation 35.1-6 on page 679. The expression is identical to the formula for the number of Lyndon words (relation 17.2-2 on page 361). If n is prime then $I_n = \frac{2^n-2}{n}$. Figure 38.6-A for gives I_n $n \le 40$, the sequence is entry A001037 of [240]. The list of all irreducible polynomials up to degree 11 is given in [FXT: data/all-irredpoly.txt].

For large degrees n the probability that a randomly chosen polynomial is irreducible is about $1/n$. With polynomials in two or more variables the situation is very different: the probability that a random polynomial is irreducible tends to one for large n, see [48].

The number of primitive binary polynomials of degree n equals

$$P_n \;=\; \frac{\varphi(2^n-1)}{n} \tag{38.6-2}$$

If n is the exponent of a Mersenne prime we have $P_n = \frac{2^n-2}{n} = I_n$. The values of P_n for $n \le 40$ are shown in figure 38.6-B. The sequence is entry A011260 of [240]. The list of all irreducible polynomials up to degree 11 is given in [FXT: data/all-primpoly.txt].

n:	D_n	n:	D_n	n:	D_n	n:	D_n
1:	1	11:	10	21:	15186	31:	0
2:	0	12:	191	22:	70525	32:	67106816
3:	0	13:	0	23:	7762	33:	49284730
4:	1	14:	405	24:	422390	34:	168436515
5:	0	15:	382	25:	46176	35:	52431606
6:	3	16:	2032	26:	860895	36:	1183272848
7:	0	17:	0	27:	768512	37:	16657254
8:	14	18:	6756	28:	4844763	38:	2411232705
9:	8	19:	0	29:	104982	39:	2168255670
10:	39	20:	28377	30:	17970267	40:	15645204474

Figure 38.6-C: The number of irreducible non-primitive binary polynomials for degrees $n \le 40$.

The difference $D_n := I_n - P_n$ is the number of irreducible non-primitive polynomials (see figure 38.6-C). Whenever n is the exponent of a Mersenne prime we have $D_n = 0$. The complete list of these polynomials up to degree 12 inclusive is given in [FXT: data/all-nonprim-irredpoly.txt].

Figure 38.6-D gives the probability that a randomly chosen irreducible polynomial of degree n is primitive. A polynomial of prime degree is very likely primitive, so any conjecture suggesting that polynomials of a certain type are always primitive for prime degree is dubious: if we take one random irreducible polynomial for each prime degree n, then chances are that all of them are primitive.

38.7 Transformations that preserve irreducibility

The *reciprocal* of a polynomial $F(x)$ is the polynomial

$$F^*(x) \;=\; x^{\deg F}\,F(1/x) \tag{38.7-1}$$

The roots of $F^*(x)$ are the inverses of the roots of $F(x)$. The reciprocal of a binary polynomial is the reversed binary word:

n :	P_n/I_n	n :	P_n/I_n	n :	P_n/I_n	n :	P_n/I_n
1:	0.50000000	26:	0.66642256	51:	0.84834222	76:	0.52983738
2:	1.0	27:	0.84540117	52:	0.51936149	77:	0.93832726
3:	1.0	28:	0.49462097	53:	0.99982834	78:	0.56391518
4:	0.66666667	29:	0.99432922	54:	0.53392943	79:	0.99962783
5:	1.0	30:	0.49790073	55:	0.91393553	80:	0.42915344
6:	0.66666667	31:	1.0	56:	0.46549716	81:	0.84506003
7:	1.0	32:	0.50000763	57:	0.85711404	82:	0.65858526
8:	0.53333333	33:	0.81066253	58:	0.65165057	83:	0.99401198
9:	0.85714286	34:	0.66665141	59:	0.99999444	84:	0.38979140
10:	0.60606061	35:	0.94659138	60:	0.35255399	85:	0.96773455
11:	0.94623656	36:	0.38011770	61:	1.0	86:	0.66505112
12:	0.42985075	37:	0.99551569	62:	0.66666667	87:	0.85207814
13:	1.0	38:	0.66666285	63:	0.83624531	88:	0.47128978
14:	0.65116279	39:	0.84618267	64:	0.49921989	89:	1.0
15:	0.82493126	40:	0.43083023	65:	0.96762379	90:	0.46446197
16:	0.50196078	41:	0.99992518	66:	0.53157031	91:	0.99091593
17:	1.0	42:	0.55199996	67:	0.99999999	92:	0.51925414
18:	0.53509496	43:	0.99757669	68:	0.52884860	93:	0.85714286
19:	1.0	44:	0.50216809	69:	0.83890107	94:	0.66388120
20:	0.45821639	45:	0.81138931	70:	0.55834947	95:	0.96267339
21:	0.84792405	46:	0.65247846	71:	0.99999560	96:	0.38730483
22:	0.62990076	47:	0.99935309	72:	0.35544000	97:	0.99991264
23:	0.97871804	48:	0.38932803	73:	0.99772166	98:	0.64603552
24:	0.39561006	49:	0.99212598	74:	0.66330362	99:	0.79553432
25:	0.96559617	50:	0.58273388	75:	0.82216371	100:	0.45025627

Figure 38.6-D: Ratios P_n/I_n for degrees $n \le 100$: a random irreducible binary polynomial of prime degree n is likely primitive even if n is not a Mersenne exponent.

```
1   inline ulong bitpol_recip(ulong c)
2   // Return x^deg(C) * C(1/x)  (the reciprocal polynomial)
3   {
4       ulong t = 0;
5       while ( c )
6       {
7           t <<= 1;
8           t |= (c & 1);
9           c >>= 1;
10      }
11      return  t;
12  }
```

Alternatively, one can use the bit-reversion routines given in section 1.14 on page 34.

The reciprocal of an irreducible polynomial is again irreducible. The order of the polynomial is preserved under the transformation.

When a polynomial is irreducible then the composition with $x + 1$ is also irreducible. Reversion also preserves primitivity, composition with $x+1$ does not in general: the most simple example is the primitive polynomial $p(x) = x^4 + x^3 + 1$ where $p(x + 1) = x^4 + x^3 + x^2 + x + 1$ has the order 5. The order of x modulo $p(x)$ equals the order of $x + 1$ modulo $p(x + 1)$.

Composition with $x + 1$ can be computed by [FXT: bpol/bitpol-irred.h]:

```
1   inline ulong bitpol_compose_xp1(ulong c)
2   // Return C(x+1).
3   // Self-inverse.
4   {
5       ulong z = 1;
6       ulong r = 0;
```

```
7       while ( c )
8       {
9           if ( c & 1 )  r ^= z;
10          c >>= 1;
11          z ^= (z<<1);
12      }
13      return  r;
14  }
```

A version that finishes in time $\log_2(b)$ (where b = bits per word) is

```
1   inline ulong bitpol_compose_xp1(ulong c)
2   {
3       ulong s = BITS_PER_LONG >> 1;
4       ulong m = ~0UL << s;
5       while ( s )
6       {
7           c ^= ( (c&m) >> s );
8           s >>= 1;
9           m ^= (m>>s);
10      }
11      return  c;
12  }
```

Which is exactly the `blue_code()` from section 1.19 on page 49.

In general the sequence of successive 'compose' and 'reverse' operations leads to 6 different polynomials:

```
C= [11, 10, 4, 3, 0]

   [11, 10, 4, 3, 0]  -- recip (C=bitpol_recip(C)) -->
   [11,  8, 7, 1, 0]   -- compose (C=bitpol_compose_xp1(C)) -->
   [11, 10, 9, 7, 6, 5, 4, 1, 0]  -- recip -->
   [11, 10, 7, 6, 5, 4, 2, 1, 0]  -- compose -->
   [11,  9, 7, 2, 0]  -- recip -->
   [11,  9, 4, 2, 0]  -- compose -->
   [11, 10, 4, 3, 0]  == initial value
```

A list of all binary irreducible polynomials of degrees $2 \le d \le 11$ is given in [FXT: data/all-irredpoly.txt].

38.8 Self-reciprocal polynomials

A polynomial is called *self-reciprocal* if it is its own reciprocal. The irreducible self-reciprocal polynomials (SRPs), except for $1+x$, are of even degree $2\,d$. They can be computed from the irreducible polynomials of degree d with nonzero linear coefficient. Let $F(x) = \sum_{j=0}^{d} f_j\, x^j$ and $S_F(x)$ the corresponding SRP, then

$$S_F(x) \;=\; x^d\, F(x+1/x) \;=\; \sum_{j=0}^{d} F_j\, x^{d-j}\, (1+x^2)^j \tag{38.8-1}$$

The irreducible SRPs of degree 18 and their corresponding polynomials are shown in figure 38.8-A [FXT: gf2n/bitpol-srp-demo.cc]. The conversion can be implemented as [FXT: bpol/bitpol-srp.h]:

```
1   inline ulong bitpol_pol2srp(ulong f, ulong d)
2   // Return  the self-reciprocal polynomial S=x^d*F(x+1/x) where d=deg(f).
3   // W = sum(j=0, d,   F(j)*x^(d-j)*(1+x^2)^j )  where
4   //   F(j) is the j-th coefficient of F.
5   // Must have: d==degree(F)
6   {
7       ulong w = 1;  // == (x^2+1)^j
8       ulong s = 0;
9       do  // for j = 0 ... d:
10      {
11          if ( f & 1 )  s ^= (w << d);  // S += F(j)*x^(d-j)*(1+x^2)^j
12          w ^= (w<<2);                  // w *= (1+x^2)
13          f >>= 1;                      // next coefficient to low end
14      }
15      while ( d-- );
```

```
              irred. poly      irred. SRP
         1:    11...1..11    111..1..111..1..111
         2:    1..111.111    1.11111.111.11111.1
         3:    1.11.11.11    1..111.11111.111..1
         4:    11111...11    11.1...1.1.1...1.11
         5:    1....1.111    1.1..1.11111.1..1.1
         6:    11.11.1.11    11111..1.1.1..11111
         7:    111...1111    11..1..11111..1..11
         8:    1....11.11    1.1..11..1..11..1.1
         9:    11.111..11    11111111111111111111
        10:    1..11.1111    1.111....1....111.1
        11:    1..1.11111    1.11.11..1..11.11.1
        12:    1111..1.11    11.11..11111..11.11
        13:    1.11..1111    1..11...111...11..1
        14:    11.1.11.11    1111.111.1.111.1111
        15:    1111111.11    11.1.111111111.1.11
        16:    1111...111    11.11.1..1..1.11.11
        17:    1.1.11.111    1....1.1.1.1.1....1
        18:    11.1..1111    1111..1..1..1..1111
        19:    11.1111111    111111...1...111111
        20:    1.1.1.1111    1.....1111111.....1
        21:    1.1.1...11    1.......1.......1
        22:    1..1..1.11    1.11..11.1.11..11.1
        23:    1.1....111    1...1.11.1.11.1...1
        24:    1.......11    1.1.....111.....1.1
        25:    11...11111    111..111.1.111..111
        26:    11..111.11    111.11...1...11.111
        27:    1...11..11    1.1.111.111.111.1.1
        28:    11..1...11    111.1.1.111.1.1.111
```

Figure 38.8-A: All irreducible self-reciprocal binary polynomials of degree 18 (right) and the corresponding irreducible polynomials of degree 9 with constant linear coefficient (left).

```
16      return s;
17  }
```

The inverse function is given in [193]:

1. Initialization: set $F := 0$, and $j := 0$.
2. If $S \bmod (x^2+1) \equiv 0$ then set $f_j := 0$, else set $f_j := 1$.
3. Set $S := (S - f_j\, x^{d-j})/(x^2+1)$ [the division is exact].
4. Set $j := j+1$. If $j \le d$ goto step 2.
5. Return F $(= \sum_{j=0}^{d} f_j\, x^j)$.

The computation of $S \bmod (x^2+1)$ can be omitted because the quantity is zero exactly if the central coefficient of S is zero. The assignment $S := (S - f_j\, x^{d-j})/(x^2+1)$ can be replaced by $S := S/(x^2+1)$ (as power series) because no coefficient beyond the position $d-j$ is needed by the following steps. We use the power series division shown in section 38.1.6 on page 820 for this computation:

```
1   inline ulong bitpol_srp2pol(ulong s, ulong hd)
2   // Inverse of bitpol_pol2srp().
3   // Must have: hd = degree(s)/2 (note: _half_ of the degree).
4   // Only the lower half coefficients are accessed, i.e.
5   //   the routine works for degree(S) <= 2*BITS_PER_LONG-2.
6   {
7       ulong f = 0;
8       ulong mh = 1UL << hd;
9       ulong ml = 1;
10      do
11      {
12          ulong b = s & mh;  // central coefficient
13  //        s ^= b;  // set central coefficient to zero (not needed)
14          if ( b )  f |= ml;  // positions 0,1,...,hd
15          ml <<= 1;
16          s = bitpol_div_x2p1(s);  // exact division by (x^2+1)
17      }
18      while ( (mh>>=1) );
19      return f;
20  }
```

The self-reciprocal polynomials of degree $2n$ are factors of the polynomial $x^{2^n+1}-1$ (see [201]). For example, for $n=5$ we obtain

```
? lift(factormod(x^(2^5+1)-1,2))
  [x + 1 1]
  [x^2 + x + 1 1]
  [x^10 + x^7 + x^5 + x^3 + 1 1]
  [x^10 + x^9 + x^5 + x + 1 1]
  [x^10 + x^9 + x^8 + x^7 + x^6 + x^5 + x^4 + x^3 + x^2 + x + 1 1]
```

The order of a self-reciprocal polynomial of degree $2n$ is a divisor of $2^{2^n}+1$. The list of all irreducible SRP up to degree 22 is given in [FXT: data/all-irred-srp.txt].

n :	S_n	n :	S_n	n :	S_n	n :	S_n
1:	1	11:	93	21:	49929	31:	34636833
2:	1	12:	170	22:	95325	32:	67108864
3:	1	13:	315	23:	182361	33:	130150493
4:	2	14:	585	24:	349520	34:	252645135
5:	3	15:	1091	25:	671088	35:	490853403
6:	5	16:	2048	26:	1290555	36:	954437120
7:	9	17:	3855	27:	2485504	37:	1857283155
8:	16	18:	7280	28:	4793490	38:	3616814565
9:	28	19:	13797	29:	9256395	39:	7048151355
10:	51	20:	26214	30:	17895679	40:	13743895344

Figure 38.8-B: Number of irreducible self-reciprocal polynomials of degree $2n$.

The number of irreducible SRPs of degree $2n$ is

$$S_n = \frac{1}{2n} \sum_{d \backslash n,\ d \text{ odd}} \mu(d)\, 2^{n/d} \tag{38.8-2}$$

Values of S_n for $n \le 40$ are shown in figure 38.8-B. The sequence of values S_n is entry A000048 of [240]. The number of irreducible polynomials of degree n where the linear coefficient is one is also S_n.

The number of irreducible SRPs of degree n is

$$\tilde{S}_n = \frac{-1}{n} \sum_{d \backslash n,\ d \text{ even}} \mu(d)\, 2^{n/d} \tag{38.8-3}$$

We have $\tilde{S}_n = 0$ for odd n and else $\tilde{S}_n = S_{n/2}$.

38.9 Irreducible and primitive polynomials of special forms *

We give lists of irreducible and primitive polynomials of special forms. For the computation of irreducible polynomials the built-in routine `polisirreducible()` of the pari/gp package was used. Primitivity was tested with the routines given so far. The abbreviation 'PP' is used for 'primitive polynomial' in what follows.

The *weight* of a binary polynomial is the sum of its coefficients.

38.9.1 All irreducible polynomials for low degrees

For degrees $n \le 8$ the complete list of irreducible polynomials is shown in figure 38.9-A. The list up to degree $n = 11$ is given in [FXT: data/all-irredpoly.txt]. The list of PPs for $n \le 11$ is given in [FXT: data/all-primpoly.txt].

```
2,1,0

3,1,0
3,2,0

4,1,0
4,3,0
# non-primitive:
4,3,2,1,0

5,2,0
5,3,0
5,3,2,1,0
5,4,2,1,0
5,4,3,1,0
5,4,3,2,0

6,1,0
6,4,3,1,0
6,5,0
6,5,2,1,0
6,5,3,2,0
6,5,4,1,0
# non-primitive:
6,3,0
6,4,2,1,0
6,5,4,2,0
```

```
7,1,0
7,3,0
7,3,2,1,0
7,4,0
7,4,3,2,0
7,5,2,1,0
7,5,3,1,0
7,5,4,3,0
7,5,4,3,2,1,0
7,6,0
7,6,3,1,0
7,6,4,1,0
7,6,4,2,0
7,6,5,2,0
7,6,5,3,2,1,0
7,6,5,4,0
7,6,5,4,2,1,0
7,6,5,4,3,2,0
```

```
8,4,3,2,0
8,5,3,1,0
8,5,3,2,0
8,6,3,2,0
8,6,4,3,2,1,0
8,6,5,1,0
8,6,5,2,0
8,6,5,3,0
8,6,5,4,0
8,7,2,1,0
8,7,3,2,0
8,7,5,3,0
8,7,6,1,0
8,7,6,3,2,1,0
8,7,6,5,2,1,0
8,7,6,5,4,2,0
# non-primitive:
8,4,3,1,0
8,5,4,3,0
8,5,4,3,2,1,0
8,6,5,4,2,1,0
8,6,5,4,3,1,0
8,7,3,1,0
8,7,4,3,2,1,0
8,7,5,1,0
8,7,5,4,0
8,7,5,4,3,2,0
8,7,6,4,2,1,0
8,7,6,4,3,2,0
8,7,6,5,4,1,0
8,7,6,5,4,3,0
```

Figure 38.9-A: All binary irreducible polynomials up to degree 8.

38.9.2 All irreducible and primitive trinomials for low degrees

```
 2,1     10,3     17,3     22,1     -30,1     35,2     -42,7
 3,1     10,7     17,5     22,21    -30,9     35,33    -42,35
 3,2     11,2     17,6     23,5     -30,21   -36,9     -44,5
 4,1     11,9     17,11    23,9     -30,29    36,11    -44,39
 4,3    -12,3     17,12    23,14     31,3    -36,15    -46,1
 5,2    -12,5     17,14    23,18     31,6    -36,21    -46,45
 5,3    -12,7    -18,3     25,3      31,7     36,25     47,5
 6,1    -12,9     18,7     25,7      31,13   -36,27     47,14
-6,3    -14,5    -18,9     25,18     31,18    39,4     47,20
 6,5    -14,9     18,11    25,22     31,24    39,8     47,21
 7,1     15,1    -18,15   -28,1      31,25    39,14    47,26
 7,3     15,4     20,3     28,3      31,28    39,25    47,27
 7,4     15,7    -20,5     28,9     -33,10    39,31    47,33
 7,6     15,8    -20,15    28,13     33,13    39,35    47,42
-9,1     15,11    20,17    28,15     33,20    41,3     49,9
 9,4     15,14    21,2     28,19    -33,23    41,20    49,12
 9,5              -21,7    28,25    -34,7     41,21    49,15
-9,8              -21,14  -28,27    -34,27    41,38    49,22
                   21,19    29,2                        49,27
                            29,27                       49,34
                                                        49,37
                                                        49,40
```

Figure 38.9-B: All irreducible trinomials $x^n + x^k + 1$ for degrees $n \leq 49$. The format of the entries is `n,k` for primitive trinomials, and `-n,k` for non-primitive trinomials.

A *trinomial* is a polynomial with exactly three nonzero coefficients. The irreducible binary trinomials for degrees $n \leq 49$ are shown in figure 38.9-B (there are no irreducible trinomials for degrees 50 and 51). A list of all irreducible trinomials up to degree $n = 400$ is given in [FXT: data/all-trinomial-irredpoly.txt]. A more compact form of the list can is given in [FXT: data/all-trinomial-primpoly-short.txt]:

```
 2:  1
 3:  1 2
 4:  1 3
 5:  2 3
 6:  1 5
 7:  1 3 4 6
 9:  4 5
10:  3 7
11:  2 9
15:  1 4 7 8 11 14
```

A line starts with the entry for the degree followed by all possible positions of the middle coefficient. The corresponding files giving primitive trinomials only are [FXT: data/all-trinomial-primpoly.txt] and [FXT: data/all-trinomial-primpoly-short.txt]. A list of irreducible trinomials that are *not* primitive is [FXT: data/all-trinomial-nonprimpoly.txt].

The sequence of values n such that a irreducible trinomial of degree n exists is entry A073571 of [240], the values for which at least one trinomial is primitive is entry A073726. The values n, k for primitive polynomials of the form $(x+1)^n + (x+1)^k + 1$ are listed in [FXT: data/all-t1-primpoly.txt]. Polynomials of that form are irreducible whenever $x^n + x^k + 1$ is. The list is not the same as for primitive trinomials as the transformation $p(x) \mapsto p(x+1)$ does in general not preserve the order. The sequence of degrees n such that there is a primitive polynomial $(x+1)^n + (x+1)^k + 1$ where $0 < k < n$ is entry A136416 of [240].

Regarding trinomials, there is a theorem by Swan (given in [249]): The trinomial $x^n + x^k + 1$ over GF(2) has an even number of irreducible factors (and so is reducible) if

1. n is even, k is odd, $n \neq 2k$, and either $nk/2 \equiv 0 \bmod 4$ or $nk/2 \equiv 1 \bmod 4$
2. n is odd, k is even and does not divide $2n$, and $n \equiv \pm 3 \bmod 8$
3. n is even, k is odd and does divide $2n$, and $n \equiv \pm 1 \bmod 8$
4. Any of the above holds for k replaced by $n - k$ (that is, for the reciprocal trinomial)

The first condition implies that no irreducible trinomial for n a multiple of 8 exists (as n is even, k must be odd, else the trinomial is a perfect square; and $nk/2 \equiv 0 \bmod 4$). Further, if n is a prime with $n \equiv \pm 3 \bmod 8$ then the trinomial can be irreducible only if $k = 2$ (or $n - k = 2$). In the note [88] it is shown that no irreducible trinomial exists for n a prime such that $n \equiv 13 \bmod 24$ or $n \equiv 19 \bmod 24$.

For some applications one may want to use reducible trinomials whose period is close to that of a primitive one. For example, the trinomial (given in [89])

$$x^{32} + x^{15} + 1 = \tag{38.9-1}$$
$$\left(x^{11} + x^9 + x^7 + x^2 + 1\right) \cdot \left(x^{21} + x^{19} + x^{15} + x^{13} + x^{12} + x^{10} + x^9 + x^8 + x^7 + x^6 + x^4 + x^2 + 1\right)$$

has the period $p = 4,292,868,097$ which is very close to $2^{32} - 1 = 4,294,967,295$. Note that the degree is a multiple of eight, so no irreducible trinomial of that degree exists.

38.9.3 Irreducible trinomials of the form $1 + \mathbf{x^k} + \mathbf{x^d}$

k=1: The polynomials of the form $1 + x + x^d$ are irreducible for the following $2 \le d \le 34353$:

```
2, 3, 4, 6, 7, 9, 15, 22, 28, 30, 46, 60, 63,
127, 153, 172, 303, 471, 532, 865, 900,
1366, 2380, 3310, 4495, 6321, 7447,
10198, 11425, 21846, 24369, 27286, 28713, 32767, 34353
```

This is sequence A002475 of [240]. The primitive trinomials of the form $x^n + x + 1$ for $n \le 400$ are those with $n \in$

```
2, 3, 4, 6, 7, 15, 22, 60, 63, 127, 153
```

These numbers are the sequence A073639 of [240], where one finds in addition 471, 532, 865, 900, 1366 with the next candidate being 4495.

k=2: $1 + x^2 + x^d$ is irreducible for the following $3 \le d \le 57341$ (sequence A057460):

```
3, 5, 11, 21, 29, 35, 93, 123, 333, 845, 4125,
10437, 10469, 14211, 20307, 34115, 47283, 50621, 57341
```

k=3: $p = 1 + x^3 + x^d$ is irreducible for the following $4 \le d \le 1000$ (sequence A057461):

```
4, 5, 6, 7, 10, 12, 17, 18, 20, 25, 28, 31, 41, 52, 66,
130, 151, 180, 196, 503, 650, 761, 986
```

k=4: $p = 1 + x^4 + x^d$ is irreducible for the following $5 \le d \le 1000$ (sequence A057463):

```
7, 9, 15, 39, 57, 81, 105
```

k=5: $p = 1 + x^5 + x^d$ is irreducible for the following $6 \le d \le 1000$ (sequence A057474):

```
6, 9, 12, 14, 17, 20, 23, 44, 47, 63, 84,
129, 236, 278, 279, 297, 300, 647, 726, 737,
```

38.9.4 Irreducible trinomials of the form $1 + x^d + x^{kd}$

k=2: The polynomial $p = 1 + x^d + x^{2d}$ is irreducible whenever d is a power of three:

```
  1:  x^2 +  x +  1
  3:  x^6 +  x^3 +  1
  9:  x^18 +  x^9 +  1
 27:  x^54 +  x^27 +  1
 81:  x^162 +  x^81 +  1
243:  x^486 +  x^243 +  1
...
```

k=3: Similarly, $p = 1 + x^d + x^{3d}$ is irreducible whenever d is a power of seven:

```
  1:  x^3 +  x +  1
  7:  x^21 +  x^7 +  1
 49:  x^147 +  x^49 +  1
343:  x^1029 +  x^343 +  1
```

k=4: The polynomial $p = 1 + x^d + x^{4d}$ whenever $d = 3^i\, 5^j$, $i, j \in \mathbb{N}$:

```
4, 12, 20, 36, 60, 100, 108, 180, 300, 324, 500, 540, 900, 972,
1500, 1620, 2500, 2700, 2916, 4500, 4860, 7500, 8100, 8748, 12500, ...
```

Similar regularities can be observed for related forms, see [45].

38.9.5 Primitive pentanomials

A *pentanomial* is a polynomial that has exactly five nonzero coefficients. PPs that are pentanomials are given in [FXT: data/pentanomial-primpoly.txt]. No primitive pentanomial exists for degrees $n < 5$ but for all higher degrees one seems to exist but this has not been proven so far. Entries with the special form $x^n + x^3 + x^2 + x + 1$ are $n \in \{5, 7, 17, 25, 31, 41, 151\}$.

38.9.6 Primitive minimum-weight and low-bit polynomials

```
[1,0]                 17,3,0                33,13,0               49,9,0
2,1,0                 18,7,0                34,8,4,3,0            50,4,3,2,0
3,1,0                 19,5,2,1,0            35,2,0                51,6,3,1,0
4,1,0                 20,3,0                36,11,0               52,3,0
5,2,0                 21,2,0                37,6,4,1,0            53,6,2,1,0
6,1,0                 22,1,0                38,6,5,1,0            54,8,6,3,0
7,1,0                 23,5,0                39,4,0                55,24,0
8,4,3,2,0             24,4,3,1,0            40,5,4,3,0            56,7,4,2,0
9,4,0                 25,3,0                41,3,0                57,7,0
10,3,0                26,6,2,1,0            42,7,4,3,0            58,19,0
11,2,0                27,5,2,1,0            43,6,4,3,0            59,7,4,2,0
12,6,4,1,0            28,3,0                44,6,5,2,0            60,1,0
13,4,3,1,0            29,2,0                45,4,3,1,0            61,5,2,1,0
14,5,3,1,0            30,6,4,1,0            46,8,7,6,0            62,6,5,3,0
15,1,0                31,3,0                47,5,0                63,1,0
16,5,3,2,0            32,7,6,2,0            48,9,7,4,0            64,4,3,1,0
```

Figure 38.9-C: Binary primitive polynomials of minimum weight.

The data in [FXT: data/minweight-primpoly.txt] lists minimal-weight PPs where in addition the coefficients are as close to the low end as possible. The first entries are shown in figure 38.9-C A list of minimal-weight PPs that fit into a machine word is given in [FXT: bpol/primpoly-minweight.cc].

Choosing those PPs where the highest nonzero coefficient is as low as possible one obtains the list in [FXT: data/lowbit-primpoly.txt]. It starts as shown in figure 38.9-D. The corresponding extract for small

[1,0]	17,3,0	33,6,4,1,0	49,6,5,4,0
2,1,0	18,5,2,1,0	34,7,6,5,2,1,0	50,4,3,2,0
3,1,0	19,5,2,1,0	35,2,0	51,6,3,1,0
4,1,0	20,3,0	36,6,5,4,2,1,0	52,3,0
5,2,0	21,2,0	37,5,4,3,2,1,0	53,6,2,1,0
6,1,0	22,1,0	38,6,5,1,0	54,6,5,4,3,2,0
7,1,0	23,5,0	39,4,0	55,6,2,1,0
8,4,3,2,0	24,4,3,1,0	40,5,4,3,0	56,7,4,2,0
9,4,0	25,3,0	41,3,0	57,5,3,2,0
10,3,0	26,6,2,1,0	42,5,4,3,2,1,0	58,6,5,1,0
11,2,0	27,5,2,1,0	43,6,4,3,0	59,6,5,4,3,1,0
12,6,4,1,0	28,3,0	44,6,5,2,0	60,1,0
13,4,3,1,0	29,2,0	45,4,3,1,0	61,5,2,1,0
14,5,3,1,0	30,6,4,1,0	46,8,5,3,2,1,0	62,6,5,3,0
15,1,0	31,3,0	47,5,0	63,1,0
16,5,3,2,0	32,7,5,3,2,1,0	48,7,5,4,2,1,0	64,4,3,1,0

Figure 38.9-D: Binary primitive polynomials of small weight and nonzero coefficient at low indices.

degrees is given in [FXT: bpol/primpoly-lowbit.cc]. The index (position) of the second highest nonzero coefficient (the *subdegree* of the polynomial) grows slowly with n and is ≤ 12 for all $n \leq 400$. Thereby one can store the list compactly as an array of 16-bit words.

38.9.7 All primitive low-bit polynomials for certain degrees

A list of all PPs $x^n + \sum_{j=0}^{k} c_j\, x^j$ for degree $n = 256$ with the second-highest order $k \leq 15$ (and the first few polynomials for $k = 16$) is given in [FXT: data/lowbit256-primpoly.txt]. The list starts as

```
256,10,5,2,0
256,10,8,5,4,1,0
256,10,9,8,7,4,2,1,0
256,11,8,4,3,2,0
256,11,8,6,4,3,0
256,11,10,9,4,2,0
256,11,10,9,7,4,0
256,12,7,5,4,2,0
256,12,8,7,6,3,0
```

Equivalent tables for degrees `DEG=` 63, 64, 127, 128, 256, 512, 521, 607, 1000, and 1024, can be found in the files `data/lowbitDEG-primpoly.txt` (where `DEG` has to be replaced by the number).

38.9.8 Primitive low-block polynomials

A low-block polynomial has the special form $x^n + \sum_{j=0}^{k} x^j$. Such PPs exist for 218 degrees $n \leq 400$. These are especially easy to store in an array (saving the index of the second highest nonzero coefficient in array element n). Moreover, simplified and possible more efficient routines based on the special forms might be possible. A complete list of all low-block PPs with degree $n \leq 400$ is given in [FXT: data/all-lowblock-primpoly.txt]. A short form of the list is [FXT: data/all-lowblock-primpoly-short.txt]. Among the low-block PPs are a few where just one bit (the coefficient after the leading coefficient) is not set. For $n \leq 400$ this is for the following degrees:

```
3, 5, 7, 13, 15, 23, 37, 47, 85, 127, 183, 365, 383
```

The PPs listed in [FXT: data/lowblock-primpoly.txt] have the smallest possible block of set bits.

38.9.9 Irreducible all-ones polynomials

Irreducible polynomials of the form $x^n + x^{n-1} + x^{n-2} + \ldots + x + 1$ (so-called *all-ones polynomials*) exist whenever $s := n + 1$ is a prime number for which 2 is a primitive root. The all-ones polynomials $x^{s-1} + \ldots + 1$ are irreducible for the following $s < 400$:

```
2, 3, 5, 11, 13, 19, 29, 37, 53, 59, 61, 67, 83, 101, 107, 131, 139, 149,
163, 173, 179, 181, 197, 211, 227, 269, 293, 317, 347, 349, 373, 379, 389
```

The sequence is entry A001122 of [240]. A list of values up to 2000 is shown in figure 39.6-B on page 866.

With the exception of $x^2 + x + 1$, none of the all-ones polynomials is primitive. In fact, the order of x equals s, which is immediate when printing the powers of x (example using $s = 5$, $p = x^4+x^3+x^2+x+1$):

```
k     x^k
0    ...1
1    ..i1
2    .i1.
3    i1..
4    1111
5    ...1  == x^5 == 1
```

The all-ones polynomials are a special case for the factorization of cyclotomic polynomials, see section 38.11 on page 850. Irreducible polynomials of high weight are considered in [6] where irreducible polynomials of the form $(x^{n+1} + 1)/(x + 1) + x^k$ up to degree 340 are given.

38.9.10 Irreducible self-reciprocal polynomials

```
2,1,0    % 1

4,3,2,1,0    % 1

6,3,0    % 1

8,5,4,3,0    % 1
8,7,6,4,2,1,0    % 1

10,7,5,3,0    % 1
10,9,5,1,0    % 1
10,9,8,7,6,5,4,3,2,1,0    % 3

12,10,7,6,5,2,0    % 1
12,10,9,8,6,4,3,2,0    % 1
12,11,10,9,8,7,6,5,4,3,2,1,0    % 5
12,11,9,7,6,5,3,1,0    % 1
12,8,7,6,5,4,0    % 1

14,9,7,5,0    % 1
14,10,8,7,6,4,0    % 1
14,11,10,9,8,7,6,5,4,3,0    % 3
14,12,10,7,4,2,0    % 3
14,12,9,8,7,6,5,2,0    % 1
14,13,10,8,7,6,4,1,0    % 1
14,13,11,7,3,1,0    % 3
14,13,12,11,10,9,7,5,4,3,2,1,0    % 1
14,13,12,9,8,7,6,5,2,1,0    % 1

16,15,8,1,0    % 1
16,12,11,8,5,4,0    % 1
16,13,12,10,8,6,4,3,0    % 1
16,13,8,3,0    % 1
16,14,12,11,8,5,4,2,0    % 1
16,14,13,11,10,9,8,7,6,5,3,2,0    % 1
16,14,13,12,10,8,6,4,3,2,0    % 1
16,14,13,12,11,9,8,7,5,4,3,2,0    % 1
16,15,13,11,10,8,6,5,3,1,0    % 1
16,15,13,12,10,9,8,7,6,4,3,1,0    % 1
16,15,13,9,8,7,3,1,0    % 1
16,15,14,12,10,8,6,4,2,1,0    % 1
16,15,14,13,11,10,8,6,5,3,2,1,0    % 1
16,15,14,13,12,11,8,5,4,3,2,1,0    % 1
16,15,14,13,9,8,7,3,2,1,0    % 1
16,15,14,8,2,1,0    % 1
```

Figure 38.9-E: Irreducible self-reciprocal polynomials up to degree 16.

A list of all irreducible self-reciprocal polynomials (see section 38.8 on page 840) up to degree 22 is given in [FXT: data/all-irred-srp.txt]. These polynomials have even degree and none of them (with the exception of $x^2 + x + 1$) is primitive. The number after the percent sign with each entry in figure 38.9-E equals $(2^{2^{n/2}} + 1)/r$ where r is the order of the polynomial with degree n.

38.9.11 Irreducible normal polynomials

The *normal* irreducible polynomials are those whose roots are linearly independent (see section 40.6 on page 890). A complete list up to degree $n = 13$ is given in [FXT: data/all-normalpoly.txt], figure 38.9-F shows the polynomials up to degree $n = 8$. Normal polynomials must have subdegree $n - 1$, that is, they are of the form $x^n + x^{n-1} + \ldots$. The condition is necessary but not sufficient: not all irreducible polynomials of subdegree $n-1$ are normal. A list of primitive normal polynomials $x^n+x^{n-1}+\ldots+x^w+1$ with w as big as possible is given in [FXT: data/highbit-normalpoly.txt]. Primitive normal polynomials

```
 2,1,0                 8,7,2,1,0              -9,8,0
                      -8,7,3,1,0               9,8,4,1,0
 3,2,0                 8,7,3,2,0               9,8,4,2,0
                      -8,7,4,3,2,1,0           9,8,4,3,2,1,0
 4,3,0                -8,7,5,1,0               9,8,5,4,0
-4,3,2,1,0             8,7,5,3,0               9,8,5,4,3,1,0
                      -8,7,5,4,0              -9,8,6,3,0
 5,4,2,1,0            -8,7,5,4,3,2,0           9,8,6,3,2,1,0
 5,4,3,1,0             8,7,6,1,0               9,8,6,4,3,1,0
 5,4,3,2,0             8,7,6,3,2,1,0           9,8,6,5,3,1,0
                      -8,7,6,4,2,1,0           9,8,6,5,3,2,0
 6,5,0                -8,7,6,4,3,2,0           9,8,6,5,4,1,0
 6,5,2,1,0             8,7,6,5,2,1,0           9,8,7,2,0
 6,5,4,1,0            -8,7,6,5,4,1,0           9,8,7,3,2,1,0
-6,5,4,2,0             8,7,6,5,4,2,0           9,8,7,5,4,3,0
                      -8,7,6,5,4,3,0           9,8,7,6,2,1,0
 7,6,0                                         9,8,7,6,3,1,0
 7,6,3,1,0                                     9,8,7,6,4,2,0
 7,6,4,1,0                                     9,8,7,6,4,3,0
 7,6,4,2,0                                     9,8,7,6,5,1,0
 7,6,5,2,0                                     9,8,7,6,5,4,3,1,0
 7,6,5,3,2,1,0
 7,6,5,4,2,1,0
```

Figure 38.9-F: All normal irreducible polynomials up to degree $n = 9$. Polynomials that are not primitive are marked with a '-'.

$x^n+x^{n-1}+x^w+\ldots+1$ where w is as small as possible are given in [FXT: data/lowbit-normalprimpoly.txt]. The all-ones polynomials are normal.

38.9.12 Irreducible alternating polynomials

The 'alternating' polynomial $1+\sum_{k=0}^{d} x^{2\,k+1} = 1+x+x^3+x^5\ldots+x^{2d+1}$ can be irreducible only if d is odd:

```
d:   (irred. poly.)
1:    x^3 + x + 1
3:    x^7 + x^5 + x^3 + x + 1
5:    x^11 + x^9 + x^7 + x^5 + x^3 + x + 1
```

The list up to $d = 1000$ (sequence A107220 of [240])

```
1, 3, 5, 7, 9, 13, 23, 27, 31, 37, 63, 69, 117, 119, 173, 219, 223,
247, 307, 363, 383, 495, 695, 987,
```

can be obtained (within about 10 minutes) via

```
for(d=1,1000, p=(1+sum(t=0,d,x^(2*t+1))); if(polisirreducible(Mod(1,2)*p),print1(d,", ")))
```

A computational simplification with modular reduction can be observed by noting that

$$\left(1+x+x^3+x^5+\ldots+x^n\right)\cdot\left(1+x^2\right) \;=\; 1+x+x^2+x^{n+2} \tag{38.9-2}$$

One does all computations modulo the product (with cheap reductions) and only reduces the final result modulo the alternating polynomial.

38.9.13 Primitive polynomials with uniformly distributed coefficients

Primitive polynomials with (roughly) equally spaced coefficients are given in [219] for degrees from 9 to 660. Polynomials with weight 5 (pentanomials) are given in [FXT: data/eq-primpoly-w5.txt], the polynomials around degree 500 are

```
498 372 247 124 0
499 380 253 125 0
500 378 250 127 0
501 375 255 125 0
502 370 240 121 0
```

The polynomials with weight 7 are given in [FXT: data/eq-primpoly-w7.txt], the list for weight 9 is [FXT: data/eq-primpoly-w9.txt].

38.10 Generating irreducible polynomials from Lyndon words

That the number of length-n Lyndon words (see section 17.2 on page 361) is equal to the number of degree-n irreducible polynomials is not a coincidence. Indeed, [79] gives an algorithm that, given a primitive polynomial, generates an irreducible polynomial from a Lyndon word: Let b be a Lyndon word, c be a irreducible polynomial of degree n and a an element of maximal order modulo c. Set $e = a^b$ and compute the polynomial $p(x)$ over $\mathrm{GF}(2^n)$, defined as

$$p(x) \;:=\; \big(x-e\big)\big(x-e^2\big)\big(x-e^4\big)\big(x-e^8\big)\cdots\big(x-e^{2^{n-1}}\big) \tag{38.10-1}$$

Then all coefficients of $p(x)$ are either zero or one and the polynomial is irreducible over $\mathrm{GF}(2)$.

```
      b            e        p                    b            e        p
   ....1 4 m    ...1.    .11..1 P             ....1 4 m    ...1.    .11..1 P
   ...1. 4      ..1..    .11..1 P             ...11 4 m    .1...    .11111
   ...11 4 m    .1...    .11111               ..1.1 2 m    .1.11    .1.1.1 red.
   ..1.. 4      .1..1    .11..1 P             ..111 4 m    ..111    .1..11 P
   ..1.1 2 m    .1.11    .1.1.1 red.          .1111 1 m    ....1    .1...1 red.
   ..11. 4      .1111    .11111
   ..111 4 m    ..111    .1..11 P
   .1... 4      .111.    .11..1 P
   .1..1 4      ..1.1    .11111
   .1.1. 2      .1.1.    .1.1.1 red.
   .1.11 4      .11.1    .1..11 P
   .11.. 4      ...11    .11111
   .11.1 4      ..11.    .1..11 P
   .111. 4      .11..    .1..11 P
   .1111 1 m    ....1    .1...1 red.
```

Figure 38.10-A: Characteristic polynomials of the powers $e = x^b$ of the primitive element x modulo $c = x^4 + x^3 + 1$ (left). If only necklaces are used as exponents b each polynomial is obtained only once (right). Irreducible polynomials are obtained for aperiodic necklaces.

An implementation in C++ is given in [FXT: `class necklace2bitpol` in bpol/necklace2bitpol.h]:

```
1   class necklace2bitpol
2   {
3   public:
4       ulong p_[BITS_PER_LONG+1];  // polynomial over GF(2**n_)
5       ulong n_;  // degree of c_
6       ulong c_;  // modulus (irreducible polynomial)
7       ulong h_;  // mask used for computation
8       ulong a_;  // generator modulo c
9       ulong e_;  // a^b
10
11  public:
12      necklace2bitpol(ulong n, ulong c=0, ulong a=0)
13          : n_(n), c_(c), a_(a)
14      {
15          if ( 0==c )  c_ = lowbit_primpoly[n];
16          if ( 0==a )  a_ = 2UL;  // 'x'
17          h_ = (highest_one(c_) >> 1);
18      }
19
20      ~necklace2bitpol()  { ; }
21
22      ulong poly(ulong b)
23      {
24          for (ulong k=0; k<=n_; ++k)  p_[k] = 0;
25          p_[0] = 1;
26          ulong e = bitpolmod_power(a_, b, c_, h_);
27          e_ = e;  // for reference
28
29          for (ulong d=1; d<=n_; ++d)
30          {
31              for (ulong k=d; k; --k)  p_[k] = p_[k-1];
32              p_[0] = 0;
33              for (ulong k=0; k<d; ++k)
34              {
35                  p_[k] ^= bitpolmod_mult( p_[k+1], e, c_, h_);
```

```
36              }
37
38              e = bitpolmod_square(e, c_, h_);
39          }
40
41          ulong p2 = 0;
42          for (ulong j=0; j<=n_; ++j)  p2 |= (p_[j] << j);
43          return p2;
44      }
45  };
```

Figure 38.10-A (left) shows all polynomials that are generated with $c = x^4 + x^3 + 1$ and the generator $a = x$. This is the output of [FXT: gf2n/necklace2irred-demo.cc]. The columns are: b and its cyclic period (an '`m`' appended if the word is the cyclic minimum), e and p where a '`P`' indicates that p is primitive. Observe that cyclic shifts of the same word give identical polynomials p. Further, if the period is not maximal then p is reducible. Restricting our attention to the necklaces b we obtain each polynomial just once (right of figure 38.10-A). The Lyndon words b give all degree-n irreducible polynomials. The primitive polynomials are exactly those where $\gcd(b,\, 2^n - 1) = 1$.

degree:	necklaces	search
24:	49 k/sec	147 k/sec
35:	17 k/sec	64 k/sec
45:	8 k/sec	36 k/sec
63:	3 k/sec	18 k/sec

Figure 38.10-B: Rate of generation of irreducible polynomials via necklaces and with exhaustive search.

To generate all irreducible binary polynomials of fixed degree use [FXT: `class all_irredpoly` in bpol/all-irredpoly.h]. The usage is shown in [FXT: gf2n/all-irredpoly-demo.cc]. For $n = 24$ one obtains $I_{24} = 698,870$ irreducible polynomials (among them are $P_{24} = 276,480$ primitive polynomials). The computation takes about 15 seconds. It turns out that the generation via exhaustive search [FXT: gf2n/bitpol-search-irred-demo.cc] is actually faster, figure 38.10-B gives the rates of generation for various degrees and both methods.

38.11 Irreducible and cyclotomic polynomials *

The primitive binary polynomials of degree n can be obtained by factoring the cyclotomic polynomial (see section 35.1.1 on page 677) Y_N over GF(2) where $N = 2^n - 1$. For example, with $n = 6$,

```
? n=6; N=2^n-1; lift( factormod(polcyclo(N),2) )
    [x^6 + x + 1  1]
    [x^6 + x^4 + x^3 + x + 1  1]
    [x^6 + x^5 + 1  1]
    [x^6 + x^5 + x^2 + x + 1  1]
    [x^6 + x^5 + x^3 + x^2 + 1  1]
    [x^6 + x^5 + x^4 + x + 1  1]
```

We use a routine (`pcfprint(N)`) that prints the N-th cyclotomic polynomial and its factors in symbolic form. With $n = 6$, $N = 2^n - 1 = 63$ we obtain

```
? n=6; N=2^n-1;
? pcfprint(N)
    63: [ 36 33 27 24 18 12 9 3 0 ]
     [ 6 1 0 ]
     [ 6 4 3 1 0 ]
     [ 6 5 0 ]
     [ 6 5 2 1 0 ]
     [ 6 5 3 2 0 ]
     [ 6 5 4 1 0 ]
```

The irreducible but non-primitive binary polynomials are factors of cyclotomic polynomials Y_d where $d \backslash N$, $d < N$ and the order of 2 modulo d equals n:

```
? fordiv(N,d,if(n==znorder(Mod(2,d)) && (d<N), pcfprint(d) ));
     9: [ 6 3 0 ]
```

```
      [ 6 3 0 ]
21:   [ 12 11 9 8 6 4 3 1 0 ]
      [ 6 4 2 1 0 ]
      [ 6 5 4 2 0 ]
```

The number of factors of Y_d equals $\varphi(d)/n$ so we can count how many degree-n irreducible polynomials correspond to which divisor of $N = 2^n - 1$:

```
 1:  [1:1]    1
 2:  [3:1]    1
 3:  [7:2]    2
 4:  [5:1]  [15:2]    3
 5:  [31:6]    6
 6:  [9:1]  [21:2]  [63:6]    9
 7:  [127:18]    18
 8:  [17:2]  [51:4]  [85:8]  [255:16]    30
 9:  [73:8]  [511:48]    56
10:  [11:1]  [33:2]  [93:6]  [341:30]  [1023:60]    99
11:  [23:2]  [89:8]  [2047:176]    186
```

Line 6 tells us that one irreducible polynomial of degree 6 is due to the factor 9, two are due to the factor 21 and the 6 primitive polynomials correspond to $N = 63$ itself which we have verified a moment ago. Further, the a polynomials corresponding to an entry `[d:a]` all have order d. The list was produced using

```
1   { for (n=1, 11,
2       print1(n,": ");
3       s = 0;
4       N = 2^n-1;
5       fordiv (N, d,
6           if ( n==znorder(Mod(2,d)) ,
7               a = eulerphi(d)/n;
8               print1(" [",d,":",a,"] ");
9               s += a;
10          );
11      );
12      print("   ",s);
13   ); }
```

38.12 Factorization of binary polynomials

This section treats the factorization of binary polynomials. The first part describes how to factorize polynomials that do not contain a square factor. The second part gives algorithms to detect and remove square factors. Finally, an algorithm to factorize arbitrary binary polynomials is given.

38.12.1 Factorization of squarefree polynomials

In order to factorize a polynomial that must not contain a square factor we will use *Berlekamp's Q-matrix algorithm* described in [37]. The algorithm consists of two main steps: the computation of the nullspace of a matrix, and a refinement phase that finds the distinct irreducible factors.

Let c be a binary polynomial of degree d. The Q-matrix is a $d \times d$ matrix whose n-th column can be computed as the binary polynomial $x^{2\,n} \pmod c$. The algorithm will use the nullspace of $Q - \mathrm{id}$.

The routine to compute the Q-matrix is [FXT: `setup_q_matrix()` in bpol/berlekamp.cc]:

```
1   void
2   setup_q_matrix(ulong c, ulong d, ulong *ss)
3   // Compute the Q-matrix for the degree-d polynomial c.
4   // Used in Berlekamp's factorization algorithm.
5   {
6       ulong h = 1UL << (d-1);
7       {
8           ulong x2 = 2UL;  // == 'x'
9           ulong q = 1UL;
10          x2 = bitpolmod_mult(x2, x2, c, h);
11          for (ulong k=0;  k<d;  ++k)
```

$c = x^7 + x + 1$ (irreducible):

```
      Q=          Q-id=          nullspace=
     1......    .......            1......
     ....1..    .1..1..
     .1..1..    .11.1..
     .....1.    ...1.1.
     ..1..1.    ..1.11.
     ......1    .....11
     ...1..1    ...1...
```

$c = x^7 + x^3 + x + 1 = (x+1)(x^2+x+1)(x^4+x+1)$:

```
      Q=          Q-id=          nullspace=
     1.....      .....          1.....
     ....1.1     .1..1.1        .1..1..
     .1..1.1     .11.1.1        .1.1.11
     .....1.     ...1.1.
     ..1.111     ..1..11
     ......1     .....11
     ...1.11     ...1.1.
```

$c = (1+x)^7 = x^7 + x^6 + x^5 + x^4 + x^3 + x^2 + x + 1$ (not square-free):

```
      Q=          Q-id=
     1...1..     ....1..        nullspace=
     .......     .1....         1......
     .1...1.     .11..1.
     .......     ...1...
     ..1...1     ..1.1.1
     .......     .....1.
     ...1...     ...1..1
```

Figure 38.12-A: The Q-matrices for three binary polynomials and the nullspaces of $Q - \mathrm{id}$.

```
12          {
13              ss[k] = q;
14              q = bitpolmod_mult(q, x2, c, h);
15          }
16          bitmat_transpose(ss, d, ss);
17      }
18  }
```

For the irreducible binary polynomial $c = x^7 + x + 1$ we obtain what is shown at the top of figure 38.12-A. This is the output of the program [FXT: gf2n/qmatrix-demo.cc]. The vector $n_0 = [1, 0, \ldots, 0]$ lies in the nullspace of $Q - \mathrm{id}$ for every polynomial c. For $c = x^7 + x^3 + x + 1 = (x+1)(x^2+x+1)(x^4+x+1)$ we obtain a nullspace of rank three (middle of figure 38.12-A). In fact, the rank r of the nullspace equals the number of distinct irreducible factors of c. That is, for polynomials containing a square factor we do not get the total number of factors. For example, with $c = (1+x)^7 = x^7 + x^6 + x^5 + x^4 + x^3 + x^2 + x + 1$ we obtain what is shown at the bottom of figure 38.12-A.

The vectors spanning the nullspace must be post-processed in order to obtain the irreducible factors of c. The algorithm can be described as follows: let F be a set of binary polynomials whose product equals c, the refinement step R_i proceeds as follows:
Let t be the i-th element of the nullspace. For each element $f \in F$ do the following: if the degree of f equals one, keep it in the set, else delete f from the set and add from the set $X = \{\gcd(f,t), \gcd(f,t+1)\}$ those elements whose degrees are greater or equal to one.

One starts with $F = \{c\}$ and does the refinement steps $R_0, R_1, \ldots R_{r-1}$ corresponding to the vectors of the nullspace. Afterwards the set F will contain exactly the distinct irreducible factors of c.

This is implemented in the routine [FXT: `bitpol_refine_factors()` in bpol/berlekamp.cc]:

```
1   ulong
2   bitpol_refine_factors(ulong *f, ulong nf, const ulong *nn, ulong r)
3   // Given the nullspace nn[0,...,r-1] of (Q-id)
4   // and nf factors f[0,...,nf-1] whose product equals c
5   // (typically nf=1 and f[0]==c)
6   // then get all r irreducible factors of c.
7   {
8       ulong ss[r];
9       for (ulong j=0; j<r; ++j) // for all elements t in nullspace
```

```
10      {
11          ulong t = nn[j];
12
13          // skip trivial elements in nullspace:
14          if ( bitpol_deg(t)==0 )  continue;
15
16          ulong sc = 0;
17          for (ulong b=0; b<nf; ++b)  // for all elements bv in set
18          {
19              ulong bv = f[b];
20              ulong db = bitpol_deg(bv);
21              if ( db <= 1 )  // bv cannot be reduced
22              {
23                  ss[sc++] = bv;
24              }
25              else
26              {
27                  for (ulong s=0; s<2; ++s)  // for all elements in GF(2)
28                  {
29                      ulong ti = t ^ s;
30                      ulong g = bitpol_gcd(bv, ti);
31                      if ( bitpol_deg(g) >= 1 )  ss[sc++] = g;
32                  }
33              }
34          }
35
36          nf = sc;
37          for (ulong k=0; k<nf; ++k)  f[k] = ss[k];
38
39          if ( nf>=r )  break;  // nf>r would be an error
40      }
41
42      return nf;
43  }
```

As can be seen, one can skip elements corresponding to constant polynomials. Further, as soon as the set F contains r elements, all factors are found and the algorithm terminates.

Now Berlekamp's algorithm can be implemented as

```
1   ulong
2   bitpol_factor_squarefree(ulong c, ulong *f)
3   // Fill irreducible factors of squarefree polynomial c into f[]
4   // Return number of factors.
5   {
6       ulong d = bitpol_deg(c);
7
8       if ( d<=1 )  // trivial cases:  0, 1, x, x+1
9       {
10          f[0] = c;
11          if ( 0==c )  d = 1;  // 0==0^1
12          return d;
13      }
14
15      ulong ss[d];
16      setup_q_matrix(c, d, ss);
17      bitmat_add_unit(ss, d);
18
19      ulong nn[d];
20      ulong r = bitmat_nullspace(ss, d, nn);
21
22
23      f[0] = c;
24      ulong nf = 1;
25      if ( r>1 )  nf = bitpol_refine_factors(f, nf, nn, r);
26
27      return r;
28  }
```

The algorithm for the computation of the nullspace was taken from [168], find the implementation in [FXT: bmat/bitmat-nullspace.cc].

Berlekamp's algorithm is given in [90] in a more general form: In order to factorize a polynomial with coefficients in the finite field GF(q) set up the Q-matrix with columns $x^{q\,i}$ and in the refinement step set $X = \{\gcd(f, t+0), \gcd(f, t+1), \ldots, \gcd(f, t+(q-1))\}$. The algorithm is efficient only if q is small.

38.12.2 Extracting the squarefree part of a polynomial

To test whether a polynomial c has a square factor one computes $g = \gcd(c, c')$ where c' is the derivative. If $g \neq 1$ then c has the square factor g: let $c = a \cdot b^2$, then $c' = a'\,b^2 + a\,2\,b\,b' = a'\,b^2$, so $\gcd(c, c') = b^2$. The corresponding routine for binary polynomials is given in [FXT: bpol/bitpol-squarefree.h]:

```
1   inline ulong bitpol_test_squarefree(ulong c)
2   // Return  0 if polynomial is squarefree
3   // else return square factor != 0
4   {
5       ulong d = bitpol_deriv(c);
6       if ( 0==d )  return  (1==c ? 0 : c);
7       ulong g = bitpol_gcd(c, d);
8       return  (1==g ? 0 : g);
9   }
```

The derivative of a binary polynomial can be computed easily [FXT: bpol/bitpol-deriv.h] (64-bit version):

```
1   inline ulong bitpol_deriv(ulong c)
2   // Return derived polynomial
3   {
4       c &= 0xaaaaaaaaaaaaaaaaUL;
5       return  (c>>1);
6   }
```

The coefficients at the even powers have to be deleted because derivation multiplies them with an even factor which equals zero modulo two.

If the derivative of a binary polynomial is zero, then it is a perfect square or a constant polynomial:

```
1   inline ulong bitpol_pure_square_q(ulong c)
2   // Return whether polynomial is a pure square != 1
3   {
4       if ( 1UL==c )  return 0;
5       c &= 0xaaaaaaaaaaaaaaaaUL;
6       return  (0==c);
7   }
```

The following routine returns zero if c is squarefree or equal to one. If c is has a square factor $s \neq 1$ then s is returned:

```
1   inline ulong bitpol_test_squarefree(ulong c)
2   {
3       ulong d = bitpol_deriv(c);
4       if ( 0==d )  return  (1==c ? 0 : c);
5       ulong g = bitpol_gcd(c, d);
6       return  (1==g ? 0 : g);
7   }
```

In case a polynomial is a perfect square then its square root can be computed as

```
1   inline ulong bitpol_pure_sqrt(ulong c)
2   {
3       ulong t = 0;
4       for (ulong mc=1,mt=1;  mc;  mc<<=2,mt<<=1)
5       {
6           if ( mc & c )  t |= mt;
7       }
8       return  t;
9   }
```

For the factorization algorithm for general polynomials we have to extract the product of all distinct irreducible factors (the squarefree part) from a polynomial. The routine [FXT: `bitpol_sreduce()` in bpol/bitpol-squarefree.cc] returns a polynomial where the even exponents are reduced:

```
1   ulong
2   bitpol_sreduce(ulong c)
3   {
4       ulong s = bitpol_test_squarefree(c);
5       if ( 0==s )  return c;  // c is squarefree
6
7       ulong f = bitpol_div(c, s);
8
9       do  // here s is a pure square and s>1
```

```
10      {
11          s = bitpol_pure_sqrt(s);
12      }
13      while ( bitpol_pure_square_q(s) );
14
15      ulong g = bitpol_gcd(s, f);
16      s = bitpol_div(s, g);
17      f = bitpol_mult(f, s);
18
19      return f;
20  }
```

With $c = f \cdot s^{k\,2^t}$ (k odd, the factors of f and s not necessarily distinct) the returned polynomial equals $f \cdot s^k$. Some examples:

$$\begin{aligned} a^2 &\mapsto a \\ a^4 &\mapsto a \\ a^3 = a\,a^2 &\mapsto a\,a \\ a^5 = a\,a^4 &\mapsto a\,a \\ a\,b^2 &\mapsto a\,b \\ a\,b\,b^2 &\mapsto a\,b\,b = a\,b^2 \mapsto a\,b \\ f \cdot s^{k\,2^t} &\mapsto f \cdot s^k \end{aligned}$$

To extract the squarefree part of a polynomial call the routine repeatedly until the returned polynomial equals the input:

```
1   inline ulong bitpol_make_squarefree(ulong c)
2   {
3       ulong z = c, t;
4       while ( z!=(t=bitpol_sreduce(z)) )  z = t;
5       return z;
6   }
```

The reduction routine will be called at most $\log_2(n)$ times for a polynomial of degree n: the worst case is a perfect power $p = a^{2^k-1}$ where $2^k - 1 \le n$. Observe that $(2^k - 1) = 1 + 2\,(2^{k-1} - 1)$, so the reduction routine will split p as $p = a\,s^2 \mapsto a\,s$ where $s = a^{2^{k-1}-1}$ is of the same form.

38.12.3 Factorization of arbitrary polynomials

The factorization routine for arbitrary binary polynomials [FXT: `bitpol_factor()` in bpol/bitpol-factor.cc] extracts the squarefree part f of its input c, uses Berlekamp's algorithm to factor f and updates the exponents according to the polynomial $s = c/f$:

```
1   ulong
2   bitpol_factor(ulong c, ulong *f, ulong *e)
3   // Factorize the binary polynomial c:
4   // c = \prod_{i=0}^{fct-1}{f[i]^e[i]}
5   // The number of factors (fct) is returned.
6   {
7       ulong d = bitpol_deg(c);
8       if ( d<=1 )  // trivial cases:  0, 1, x, x+1
9       {
10          f[0] = c;
11          if ( 0==c )  d = 1;  // 0==0^1
12          return d;
13      }
14
15      // get squarefree part:
16      ulong cf = bitpol_make_squarefree(c);
17
18      // ... and factor it:
19      ulong fct = bitpol_factor_squarefree(cf, f);
20
21      // All exponents are one:
22      for (ulong j=0; j<fct; ++j)  { e[j] = 1; }
```

```
23
24      // Here f[],e[] is a valid factorization of the squarefree part cf
25
26      // Update exponents with square part:
27      ulong cs = bitpol_div(c, cf);
28      for (ulong j=0; j<fct; ++j)
29      {
30          if ( 1==cs )  break;
31          ulong fj = f[j];
32          ulong g = bitpol_gcd(cs, fj);
33          while ( 1!=g )
34          {
35              ++e[j];
36              cs = bitpol_div(cs, fj);
37              if ( 1==cs )  break;
38              g = bitpol_gcd(cs, fj);
39          }
40      }
41
42      return fct;
43  }
```

As given the algorithm makes just one call to the routine that computes a nullspace.

```
 0:    x^5                   ==   (x)^5
 1:    x^5                +1 ==   (x+1) *  (x^4+x^3+x^2+x+1)
 2:    x^5             +x    ==   (x) *  (x+1)^4
 3:    x^5             +x+1  ==   (x^2+x+1) *  (x^3+x^2+1)
 4:    x^5         +x^2      ==   (x)^2 *  (x+1) *  (x^2+x+1)
 5:    x^5         +x^2  +1  ==   (x^5+x^2+1)
 6:    x^5         +x^2+x    ==   (x) *  (x^4+x+1)
 7:    x^5         +x^2+x+1  ==   (x+1)^2 *  (x^3+x+1)
 8:    x^5    +x^3           ==   (x)^3 *  (x+1)^2
 9:    x^5    +x^3       +1  ==   (x^5+x^3+1)
10:    x^5    +x^3    +x    ==   (x) *  (x^2+x+1)^2
11:    x^5    +x^3    +x+1  ==   (x+1) *  (x^4+x^3+1)
12:    x^5    +x^3+x^2      ==   (x)^2 *  (x^3+x+1)
13:    x^5    +x^3+x^2  +1  ==   (x+1)^3 *  (x^2+x+1)
14:    x^5    +x^3+x^2+x    ==   (x) *  (x+1) *  (x^3+x^2+1)
15:    x^5    +x^3+x^2+x+1  ==   (x^5+x^3+x^2+x+1)
16:    x^5+x^4              ==   (x)^4 *  (x+1)
17:    x^5+x^4           +1 ==   (x^2+x+1) *  (x^3+x+1)
18:    x^5+x^4         +x   ==   (x) *  (x^4+x^3+1)
19:    x^5+x^4         +x+1 ==   (x+1)^5
20:    x^5+x^4    +x^2      ==   (x)^2 *  (x^3+x^2+1)
21:    x^5+x^4    +x^2  +1  ==   (x+1) *  (x^4+x+1)
22:    x^5+x^4    +x^2+x    ==   (x) *  (x+1)^2 *  (x^2+x+1)
23:    x^5+x^4    +x^2+x+1  ==   (x^5+x^4+x^2+x+1)
24:    x^5+x^4+x^3          ==   (x)^3 *  (x^2+x+1)
25:    x^5+x^4+x^3       +1 ==   (x+1)^2 *  (x^3+x^2+1)
26:    x^5+x^4+x^3    +x    ==   (x) *  (x+1) *  (x^3+x+1)
27:    x^5+x^4+x^3    +x+1  ==   (x^5+x^4+x^3+x+1)
28:    x^5+x^4+x^3+x^2      ==   (x)^2 *  (x+1)^3
29:    x^5+x^4+x^3+x^2  +1  ==   (x^5+x^4+x^3+x^2+1)
30:    x^5+x^4+x^3+x^2+x    ==   (x) *  (x^4+x^3+x^2+x+1)
31:    x^5+x^4+x^3+x^2+x+1  ==   (x+1) *  (x^2+x+1)^2
```

Figure 38.12-B: Factorizations of the binary polynomials of degree 5.

Figure 38.12-B shows the factorizations of the binary polynomials of degree 5, it was created with the program [FXT: gf2n/bitpolfactor-demo.cc]. Factoring the first million polynomials of degrees 20, 30, 40 and 60 takes about 5, 10, 15 and 30 seconds, respectively.

A variant of the factorization algorithm often given uses the so-called *squarefree factorization* $c = \prod_i a_i^i$ where the polynomials a_i are squarefree and pairwise coprime. Given the squarefree factorization one has to call the core routine for each non-trivial a_i.

As noted, the refinement step becomes expensive if the coefficients are in a field GF(q) where q is not small because q computations of the polynomial gcd are involved. For an algorithm that is efficient also for large values of q see [90] or [127, ch.14].

Chapter 39

Shift registers

We describe shift register sequences (SRS) and their generation via linear feedback shift registers (LFSR). The underlying mechanism is the modular arithmetic of binary polynomials described in section 38.3. We give an expression for the number of shift registers sequences of maximal length, the so-called m-sequences. Two related mechanisms, feedback carry shift registers (FCSR) and linear hybrid cellular automata (LHCA), are described. Most of the algorithms given can easily be implemented in hardware.

39.1 Linear feedback shift registers (LFSR)

```
      c = 1..11 == 0x13 == 19  (deg = 4)
       0    w =   15  = 1111  1  ...1 =    1 = a
       1    w =   14  = 111.  .  ..1. =    2 = a
       2    w =   12  = 11..  .  .1.. =    4 = a
       3    w =    8  = 1...  .  1... =    8 = a
       4    w =    1  = ...1  1  ..11 =    3 = a
       5    w =    2  = ..1.  .  .11. =    6 = a
       6    w =    4  = .1..  .  11.. =   12 = a
       7    w =    9  = 1..1  1  1.11 =   11 = a
       8    w =    3  = ..11  1  .1.1 =    5 = a
       9    w =    6  = .11.  .  1.1. =   10 = a
      10    w =   13  = 11.1  1  .111 =    7 = a
      11    w =   10  = 1.1.  .  111. =   14 = a
      12    w =    5  = .1.1  1  1111 =   15 = a
      13    w =   11  = 1.11  1  11.1 =   13 = a
      14    w =    7  = .111  1  1..1 =    9 = a
   >>  0    w =   15  = 1111  1  ...1 =    1 = a  << period restarts
       1    w =   14  = 111.  .  ..1. =    2 = a
       2    w =   12  = 11..  .  .1.. =    4 = a
       3    w =    8  = 1...  .  1... =    8 = a
       4    w =    1  = ...1  1  ..11 =    3 = a
       5    w =    2  = ..1.  .  .11. =    6 = a
```

Figure 39.1-A: Linear feedback shift register using the primitive polynomial $C = x^4 + x + 1$.

Multiplication of a binary polynomial A by x modulo a polynomial C is particularly easy as shown near the beginning of section 38.3 on page 827: Simply shift the input to the left (multiplication) and if the result $A \cdot x$ has the same degree as C then subtract (XOR) the polynomial C (modular reduction).

The underlying mechanism of shifting and conditionally feeding back certain bits is called a *linear feedback shift register* (LFSR). A *shift register sequence* (SRS) can be obtained by computing $A_k = x^k$, $k = 0, 1, \ldots, 2^n - 1$ modulo C and setting bit k of the SRS to the least significant bit of A_k. In the context of LFSRs the polynomial C is sometimes called the *connection polynomial* of the shift register.

When the modulus C is a primitive polynomial (see section 38.5 on page 834) of degree n then the SRS is a sequence of zeros and ones that contains all nonzero words of length n. Further, if a word W is

updated at each step by left shifting and adding the bit of the SRS then this sequence also contains all nonzero words.

This is demonstrated in [FXT: gf2n/lfsr-demo.cc], which for $n = 4$ uses the primitive polynomial $C = x^4 + x + 1$ and gives the output shown in figure 39.1-A. Here we pasted the first few lines after the end of the actual output to emphasize the periodicity of the sequences. The corresponding SRS of period 15 is (extra spaces to mark start of period):

```
1 0 0 0 1 0 0 1 1 0 1 0 1 1 1    1 0 0 0 1 0 0 1 1 0 1 0 1 1 1    1 0 ...
```

In fact any of the bits of the words $A_k = x^k \bmod C$ (or linear combination of two or more bits) could be used, each producing a cyclically shifted version of the SRS.

The cheapest way to generate a SRS is to compute the negative powers x^{-k} modulo C, that is, to repeatedly divide by x:

```
1       ulong c = a primitive polynomial;
2       ulong n = degree of C;
3       ulong a = 1;
4       for (ulong k=0; k<n; ++k)
5       {
6           ulong s = a & 1;
7           // Use s here.
8           if ( s )  a ^= c;
9           a >> 1;
10      }
```

The routine will work for $\deg(C) < n$ where n is the number of bits in a machine word. A version that also works for $\deg(C) = n$ is given in section 38.3 on page 827.

A C++ class implementing a LFSR with a few convenient extras is [FXT: `class lfsr` in bpol/lfsr.h]:

```
1   class lfsr
2   // (binary) Linear Feedback Shift Register
3   // Produces a shift register sequence (SRS)
4   // generated by  a_k=x^k (mod c) where
5   // c is a primitive polynomial of degree n.
6   // The period of SRS is 2^n - 1
7   // (non-primitive c lead to smaller periods)
8   {
9   public:
10      ulong a_;  // internal state (polynomial modulo c)
11      ulong w_;  // word of the shift_register_sequence (SRS)
12      ulong c_;  // (mod 2) poly  e.g. x^4+x+1 == 0x13 == 1..11
13      ulong h_;  // highest bit in SRS word  e.g. (above) == 16 = 1...
14      ulong mask_;  // mask  e.g. (above) == 15 == 1111
15      ulong n_;  // degree of polynomial  e.g. (above) == 4
16
17  public:
18      lfsr(ulong n, ulong c=0)
19      // n: degree of polynomial c
20      // c: polynomial (defaults to minimum weight primitive polynomial)
21      {
22   [--snip--]
```

The crucial computation is implemented as

```
1       ulong next()
2       {
3           ulong s = a_ & h_;
4           a_ <<= 1;
5           w_ <<= 1;
6           if ( 0!=s )
7           {
8               a_ ^= c_;
9               w_ |= 1;
10          }
11          w_ &= mask_;
12          return w_;
13      }
```

Up to the lines that update the word `w_` this function is identical to `bitpolmod_times_x()` given in section 38.3 on page 827.

The method `next_w()` skips to the next word by calling `next()` n times:

```
1       ulong next_w()
2       {
3           for (ulong k=0; k<n_; ++k)  next();
4           return w_;
5       }
```

Let a and w a pair of values that correspond to each other. The following two methods directly set one of these two while keeping the pair consistent. Setting a to a given value:

```
1       void set_a(ulong a)
2       {
3           a_ = a;
4           w_ = 0;
5           ulong b = 1;
6           for (ulong j=0; j<n_; ++j)
7           {
8               if ( a & 1 )
9               {
10                  w_ |= b;
11                  a ^= c_;
12              }
13              b <<= 1;
14              a >>= 1;
15          }
16      }
```

The loop executes n times where n is the degree of the modulus. Setting w to a given value:

```
1       void set_w(ulong w)
2       {
3           w_ = w;
4           a_ = 0;
5           ulong c = c_;
6           while ( w )
7           {
8               if ( w & 1 )  a_ ^= c;
9               c <<= 1;
10              w >>= 1;
11          }
12          a_ &= mask_;
13      }
```

The supplied value must be nonzero for both methods.

Going back one step is possible via the method `prev()`

```
1   public:
2       ulong prev()
3       {
4           prev_a();
5           set_a(a_);
6           return w_;
7       }
```

which calls `prev_a()`:

```
1   private:
2       void prev_a()
3       {
4           ulong s = a_ & 1;
5           a_ >>= 1;
6           if ( s )
7           {
8               a_ ^= (c_>>1);
9               a_ |= h_;  // so it works for  n_ == BITS_PER_LONG
10          }
11      }
```

The method `prev_a()` leaves the value of w inconsistent with a and therefore cannot be called directly. Note that stepping back is more expensive than stepping forward because `set_a()` is rather expensive.

It is also possible to go backwards word-wise:

```
1       ulong prev_w()
2       {
3           for (ulong k=0; k<n_; ++k)  prev_a();
4           set_a(a_);
5           return w_;
6       }
```

As this routine involves only one call to `set_a()` it is about as expensive as stepping one word forward using `next_w()`.

39.2 Galois and Fibonacci setup

The type of shift registers considered so far is the so-called *Galois setup* of a binary shift register. The mechanism is to detect whether a one is being shifted out and, if so, subtract the polynomial modulus. The left- and right- shift operations can be implemented as

```
1   ulong galois_left(ulong x, ulong c, ulong h)
2   {
3       ulong s = x & h;
4       x <<= 1;
5       if ( 0!=s )  x ^= c;
6       return x;
7   }
```

and

```
1   ulong galois_right(ulong x, ulong c, ulong)
2   {
3       ulong s =  ( x & 1UL );
4       x >>= 1;
5       if ( s )  x ^= (c>>1);
6       return x;
7   }
```

```
    c = .11..1 == 0x19 == 25  (deg = 4)
    r = .1..11 == 0x13 == 19  (deg = 4)

           -------- Galois -------------          --------- Fibonacci ---------
     k:      Lc      Lr      Rc      Rr            Lc      Lr      Rc      Rr
     1:    ...1    ...1    ...1    ...1          ...1    ...1    ...1    ...1
     2:    ..1.    ..1.    .11.    .1..          ..1.    ..11    .1..    .1..
     3:    .1..    .1..    ..11    .11.          .1..    .111    .11.    ..1.
     4:    1...    1...    ...1    1111          1..1    1111    1111    ...1
     5:    1..1    ..11    .11.    111.          ..11    111.    1111    1..1
     6:    1.11    .11.    1.1.    ..11          .11.    11.1    .111    11..
     7:    1111    11..    .1.1    1.1.          11.1    1.1.    1.11    .11.
     8:    .111    1.11    111.    .1.1          1.1.    .1.1    .1.1    1.11
     9:    111.    .1.1    .111    1.11          .1.1    1.11    1.1.    .1.1
    10:    .1.1    1.1.    1111    11..          1.11    .11.    11.1    1.1.
    11:    1.1.    .111    1.11    .11.          .111    11..    .11.    11.1
    12:    11.1    111.    1..1    ...1          1111    1..1    ...1    111.
    13:    ...1    1111    1...    1...          111.    ...1    1..1    1111
    14:    .11.    11.1    .1..    .1..          11..    .1..    .1..    .111
    15:    11..    1..1    ...1    ...1          1...    1...    ...1    ...1
    16:    ...1    ...1    ...1    ...1          ...1    ...1    ...1    ...1
```

Figure 39.2-A: Sequences of words generated with the Galois- and Fibonacci- mechanisms, either with the left- or the right- shift (capital letters '`L`' and '`R`' on top of columns) and primitive polynomial '`c`' or its reciprocal '`r`'. Each track of any sequence is a shift register sequence.

Four sequences of binary words that (starting identically and) are generated with either the left- or right-shift and a primitive polynomial or its reciprocal (reversed word) are shown in figure 39.2-A. A different set of sequences shown in the same figure is obtained via the so-called *Fibonacci setup*. In the Fibonacci setup the sum (modulo 2) of bits determined by the used polynomial is shifted in at each step.

The left- and right- shift operations can be implemented as

```
1   ulong fibonacci_left(ulong x, ulong c, ulong h)
2   {
3       x <<= 1;
```

```
4       ulong s = parity( x & c );
5       if ( 0!=s )  x ^= 1;
6       x &= ~(h<<1);  // remove excess bit at high end
7       return x;
8   }
```

and

```
1   ulong fibonacci_right(ulong x, ulong c, ulong h)
2   {
3       ulong s = parity( x & c );
4       x >>= 1;
5       if ( s )  x ^= h;
6       return x;
7   }
```

As the parity computation is expensive on most machines (see section 1.16.1 on page 42) the Galois setup should usually be preferred. The programs [FXT: gf2n/lfsr-fibonacci-demo.cc] and [FXT: gf2n/lfsr-galois-demo.cc] can be used to create the binary patterns shown in figure 39.2-A. With both the polynomial modulus can be specified.

39.3 Generating all revbin pairs

```
polynomial:  c      cr                    c      cr
           1.1111  1111.1                1.1111  1111.1

      k:     x       xr            k:     x       xr
      1:   ...1    1...            17:   11...    ...11
      2:   1.111   111.1           18:   .11..    ..11.
      3:   111..   ..111           19:   ..11.    .11..
      4:   .111.   .111.           20:   ...11    11...
      5:   ..111   111..           21:   1.11.    .11.1
      6:   1.1..   ..1.1           22:   .1.11    11.1.
      7:   .1.1.   .1.1.           23:   1..1.    .1..1
      8:   ..1.1   1.1..           24:   .1..1    1..1.
      9:   1.1.1   1.1.1           25:   1..11    11..1
     10:   111.1   1.111           26:   1111.    .1111
     11:   11..1   1..11           27:   .1111    1111.
     12:   11.11   11.11           28:   1....    ....1
     13:   11.1.   .1.11           29:   .1...    ...1.
     14:   .11.1   1.11.           30:   ..1..    ..1..
     15:   1...1   1...1           31:   ...1.    .1...
     16:   11111   11111           32:   ....1    1.... <--= initial pair
```

Figure 39.3-A: All nonzero 5-bit revbin pairs generated by a LFSR.

Using a primitive polynomial of degree n and its reverse one can generate all nonzero pairs `x` and `revbin(x,n)` as follows [FXT: gf2n/lfsr-revbin-demo.cc]:

```
1   inline void revbin_next(ulong &x, ulong c, ulong &xr, ulong cr)
2   // if x and xr are (nonzero) n-bit words that are a revbin pair
3   // compute the next revbin pair.
4   // c must be a primitive polynomial, cr its reverse (the reciprocal polynomial).
5   {
6       ulong s =  ( x & 1UL );
7       x >>= 1;
8       xr <<= 1;
9       if ( s )
10      {
11          x ^= (c>>1);
12          xr ^= (cr);
13      }
14  }
```

An equivalent technique for computing the revbin permutation (see section 2.1 on page 93) has been proposed in [209]. Figure 39.3-A shows all nonzero 5-bit revbin pairs generated with the primitive polynomial $c = x^5 + x^3 + x^2 + x + 1$ and its reverse.

39.4 The number of m-sequences and De Bruijn sequences

```
degree = 2
c=111 :  .11

degree = 3
c=1.11 :  ..1.111
c=11.1 :  ..111.1

degree = 4
c=1..11 :  ...1..11.1.1111
c=11..1 :  ...1111.1.11..1

degree = 5
c=1..1.1 :  ....1..1.11..11111...11.111.1.1
c=1111.1 :  ....11..1..11111.111...1.1.11.1
c=11.111 :  ....111..11.11111.1...1..1.1.11
c=1.1111 :  ....1.11.1.1...111.11111..1..11
c=111.11 :  ....11.1.1..1...1.11111.11..111
c=1.1..1 :  ....1.1.111.11...11111..11.1..1

degree = 6
c=1...11 :  .....1...11...1.1..1111.1...111..1..1.11.111.11..11.1.1.111111
c=11..111 :  .....1111..1..1.1.1..11.1....1...1.11.111111.1.111...11..111.11
c=1.11.11 :  .....1.111111..1.1.1...11..1111.111.1.11.1..11.11...1..1....111
c=11.11.1 :  .....111....1..1...11.11..1.11.1.111.1111..11...1.1.1..111111.1
c=111..11 :  .....11.111..11...111.1.111111.11.1...1....1.11..1.1.1..1..1111
c=11....1 :  .....111111.1.1.11..11.111.11.1..1..111...1.1111..1.1...11....1
```

Figure 39.4-A: All m-sequences for $n = 2$, 3, 4, 5, and 6. Dots denote zeros.

The shift register sequences generated with a polynomial of degree n is of maximal length if the polynomial is primitive. The corresponding shift register sequences are called *m-sequences*.

We now consider all sequences that are cyclic shifts of each other as the same sequence. For given n there are as many m-sequences as primitive polynomials ($P_n = \varphi(2^n - 1)/n$, see section 38.6 on page 837). These can be generated using the linear feedback shift registers described in section 39.1 on page 857.

One might suspect that the using the powers of other elements than x might lead to additional m-sequences, but this is not the case. Further, the powers of elements of maximal order modulo irreducible non-primitive polynomials do not give additional m-sequences.

The program [FXT: gf2n/all-primpoly-srs-demo.cc] computes all m-sequences for a given n. The output for $n = 2, 3, 4, 5, 6$ is shown in figure 39.4-A.

```
      n = 4  nn = 16            n = 5  nn = 32
      ....1111.11..1.1          .....11111.111..11.1.11...1.1..1
      ....1111.1.11..1          .....11111.111..11.1.11...1..1.1
      ....1111.1..1.11          .....11111.111..11.1.1..1.11...1
      ....1111..1.11.1          .....11111.111..11.1.1..1...1.11
      ....11.1111..1.1          .....11111.111..11.1.1...1...1.11
      ....11.1.1111..1          .....11111.111..11.1..1.1.11...1
      ....11.1..1.1111          .....11111.111..11.1..1...1.1.11
      ....11..1.1111.1           [--snip--]
      ....1.1111.1..11          .....1...11..1.1.11.111.1..11111
      ....1.1111..11.1          .....1...11..1.1.11.1..11111.111
      ....1.11.1..1111          .....1...11..1.1.11.1..111.11111
      ....1.11..1111.1          .....1...11..1.1..11111.11.1.111
      ....1.1..1111.11          .....1...11..1.1..11111.1.11.111
      ....1.1..11.1111          .....1...11..1.1..111.11.1.11111
      ....1..1111.1.11          .....1...11..1.1..111.1.11.11111
      ....1..11.1.1111          total # of DBS found = 2048
total # of DBS found = 16
```

Figure 39.4-B: All De Bruijn sequences for $n = 4$ (left), the first and last sequence correspond to the two m-sequences for $n = 4$. The first and last few De Bruijn sequences for $n = 5$ (right).

When a zero is inserted to the (unique) run of $n - 1$ zeros in an m-sequence then a *De Bruijn sequence* (DBS) is obtained. A DBS contains all binary words including the all-zero word.

For all $n \geq 4$ given there exist more DBSs than m-sequences. For example, for $n = 6$ there are 6 m-sequences and 67,108,864 DBSs. An exhaustive search for all DBSs of given length $L = 2^n$ is possible only for tiny n. The program [FXT: bits/all-dbs-demo.cc] finds all DBSs for $n = 3, 4, 5$. Its output with $n = 4$ and $n = 5$ (partly) is shown in figure 39.4-B.

n :	$L_n = 2^n$	$x = 2^{n-1} - n$	$S_n = 2^x$
1:	2	0	1
2:	4	0	1
3:	8	1	2
4:	16	4	16
5:	32	11	2,048
6:	64	26	67,108,864
7:	128	57	144,115,188,075,855,872
8:	256	120	1,329,227,995,784,915,872,903,807,060,280,344,576

Figure 39.4-C: The number S_n of DBSs of length L_n.

The total number of DBSs equals

$$S_n \quad = \quad 2^x \quad \text{where} \quad x = 2^{n-1} - n \tag{39.4-1}$$

The two DBSs for $n = 3$ are `[...111.1]` and `[...1.111]=[1.111...]`, reversed sequences are considered different by the formula. The first few values of S_n are shown in figure 39.4-C. One has $S_{n+1} = S_n^2 \, L_{n-1}$, equivalently $x_{n+1} = 2\,x_n + n - 1$.

The general formula for the number of length-n base-m DBSs is $S_n = m!^{m^{n-1}}/m^n$, as given in [170]. A graph theoretical proof of the formula for the case $m = 2$ can be found in [186, p.56]. For a more efficient approach to generate all DBSs of given length see section 19.2.2 on page 377.

We note that there are several ways to generalize the idea of the De Bruijn cycles to *universal cycles* as described in [85].

39.5 Auto correlation of m-sequences

```
 1:  [ +-++--+-+----+++ ]   C= +15 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1
 2:  [ +-++----+-+--+++ ]   C= +15 -1 -1 -1 -1 +3 -5 -1 -1 -5 +3 -1 -1 -1 -1
 3:  [ +-+-++--+----+++ ]   C= +15 -1 -1 -1 -1 -1 +3 -5 -5 +3 -1 -1 -1 -1 -1
 4:  [ +-+-++----+--+++ ]   C= +15 -1 -1 -1 -1 +3 -1 -5 -5 -1 +3 -1 -1 -1 -1
 5:  [ +-+--++----+-+++ ]   C= +15 -1 -1 -1 -1 -5 +3 -1 -1 +3 -5 -1 -1 -1 -1
 6:  [ +-+--+-++----+++ ]   C= +15 -1 -1 -1 -9 -1 +3 +3 +3 +3 -1 -9 -1 -1 -1
 7:  [ +-+----++--+-+++ ]   C= +15 -1 -1 -1 -1 -5 +3 -1 -1 +3 -5 -1 -1 -1 -1
 8:  [ +-+----+-++--+++ ]   C= +15 -1 -1 -1 -1 -1 -5 +3 +3 -5 -1 -1 -1 -1 -1
 9:  [ +--++-+----+-+++ ]   C= +15 -1 -1 -1 -1 -1 -5 +3 +3 -5 -1 -1 -1 -1 -1
10:  [ +--+-++-+----+++ ]   C= +15 -1 -1 -1 -9 +3 -1 +3 +3 -1 +3 -9 -1 -1 -1
11:  [ +--+-+----++-+++ ]   C= +15 -1 -1 -1 -1 +3 -5 -1 -1 -5 +3 -1 -1 -1 -1
12:  [ +--+----++-+-+++ ]   C= +15 -1 -1 -1 -1 +3 -1 -5 -5 -1 +3 -1 -1 -1 -1
13:  [ +----++-+--+-+++ ]   C= +15 -1 -1 -1 -9 -1 +3 +3 +3 +3 -1 -9 -1 -1 -1
14:  [ +----+-++-+--+++ ]   C= +15 -1 -1 -1 -9 +3 -1 +3 +3 -1 +3 -9 -1 -1 -1
15:  [ +----+-+--++-+++ ]   C= +15 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1
16:  [ +----+--++-+-+++ ]   C= +15 -1 -1 -1 -1 -1 +3 -5 -5 +3 -1 -1 -1 -1 -1
```

Figure 39.5-A: Cyclic auto correlations for all truncated (signed) De Bruijn sequences for $n = 4$. Only two (the first and the second last) out of the 16 sequences are m-sequences.

We have seen that a De Bruijn sequence (DBS) can be obtained from an m-sequence by inserting a single zero at the longest run of zeros. In the other direction, if we take a DBS and delete a zero from the longest run of zeros then we obtain a sequence of length $N = 2^n - 1$ that contains every n-bit nonzero word. But these sequences are *not* m-sequences in general: most of them can not be obtained with an n-bit LFSR and miss an important property of m-sequences.

For a sequence M of $N-1$ zeros and ones M_k define the sequence S of elements S_k via

$$S_k \;:=\; \begin{cases} +1 & \text{if} \quad M_k = 1 \\ -1 & \text{else} \end{cases} \tag{39.5-1}$$

Then, if M is an length-L m-sequence, we have for the cyclic auto correlation (or *auto correlation function*, ACF) of S

$$C_\tau \;:=\; \sum_{k=0}^{L-1} S_k\, S_{k+\tau \bmod L} \;=\; \begin{cases} L & \text{if} \quad \tau = 0 \\ -1 & \text{else} \end{cases} \tag{39.5-2}$$

where $L = N-1$ and $N = 2^n$. That is, C_0 equals the length of the sequence, all other entries are of minimal absolute value: they cannot be zero because L is odd.

This property does not hold for most of the 'truncated' DBS (where one zero in the single run of n consecutive zeros is removed). Figure 39.5-A shows all (signed) truncated DBS for $n = 4$ and their auto correlations. Only 2 out of the 16 truncated DBS have an auto correlation satisfying relation 39.5-2, these are exactly the m-sequences for $n = 4$.

For odd primes q one can obtain sequences of length $L = q$ whose ACF satisfies

$$\sum_{k=0}^{L-1} S_k\, S_{k+\tau \bmod L} \;=\; \begin{cases} L-1 & \text{if} \quad \tau = 0 \\ -1 & \text{else} \end{cases} \tag{39.5-3}$$

However, the sequences start with a single zero: set $S_0 = 0$ and for $1 \le k < q$ set $S_k = 1$ if k is s square modulo q, else $S_k = -1$. The first three such sequences for primes of the form $4\,k+1$ and their ACFs are:

```
  5:   S: [0, +1, -1, -1, +1]
       C: [4, -1, -1, -1, -1]

 13:   S: [ 0, +1, -1, +1, +1, -1, -1, -1, -1, +1, +1, -1, +1]
       C: [12, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1]

 17:   S: [ 0, +1, +1, -1, +1, -1, -1, -1, +1, +1, -1, -1, -1, +1, -1, +1, +1]
       C: [16, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1]
```

With primes of the form $q = 4\,k+3$ one can construct a sequence of length $L = q$ satisfying relation 39.5-2 by simply setting $S_0 = 1$ ($S_0 = -1$ also works) in the sequence just constructed. The sequences for the first three primes $q = 4\,k+3$ and their ACFs are:

```
  3:   S: [+1, +1, -1]
       C: [ 3, -1, -1]

  7:   S: [+1, +1, +1, -1, +1, -1, -1]
       C: [ 7, -1, -1, -1, -1, -1, -1]

 11:   S: [+1, +1, -1, +1, +1, +1, -1, -1, -1, +1, -1]
       C: [11, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1]
```

These sequences can be used for the construction of Hadamard matrices, see section 18 on page 365.

39.6 Feedback carry shift register (FCSR)

There is a nice analogue of the LFSR in the modulo world, the *feedback carry shift register* (FCSR). With the LFSR we needed an irreducible ('prime') polynomial C where x has maximal order. The powers of x modulo C did run through all different (nonzero) words. Now take a prime c where 2 has maximal order (that is, 2 is primitive root modulo c, see section 37.4 on page 761). Then the powers of 2 modulo c run through all nonzero values less than c.

The implementation is [FXT: `class fcsr` in bpol/fcsr.h]:

```
          c = 1..1.1  = 37

          0 :   a= .....1 =    1        w= .1..11 =   19
          1 :   a= ....1. =    2        w= 1..11. =   38
          2 :   a= ...1.. =    4        w= ..11.. =   12
          3 :   a= ..1... =    8        w= .11... =   24
          4 :   a= .1.... =   16        w= 11.... =   48
          5 :   a= 1..... =   32        w= 1....., =   32
          6 :   a= .11.11 =   27        w= .....1 =    1
          7 :   a= .1...1 =   17        w= ....11 =    3
          8 :   a= 1...1. =   34        w= ...11. =    6
          9 :   a= .11111 =   31        w= ..11.1 =   13
         10 :   a= .11..1 =   25        w= .11.11 =   27
         11 :   a= ..11.1 =   13        w= 11.111 =   55
         12 :   a= .11.1. =   26        w= 1.111. =   46
         13 :   a= ..1111 =   15        w= .111.1 =   29
         14 :   a= .1111. =   30        w= 111.1. =   58
         15 :   a= .1.111 =   23        w= 11.1.1 =   53
         16 :   a= ..1..1 =    9        w= 1.1.11 =   43
         17 :   a= .1..1. =   18        w= .1.11. =   22
         18 :   a= 1..1.. =   36        w= 1.11.. =   44
         19 :   a= 1...11 =   35        w= .11..1 =   25
         20 :   a= 1....1 =   33        w= 11..11 =   51
         21 :   a= .111.1 =   29        w= 1..111 =   39
         22 :   a= .1.1.1 =   21        w= ..1111 =   15
         23 :   a= ...1.1 =    5        w= .11111 =   31
         24 :   a= ..1.1. =   10        w= 11111. =   62
         25 :   a= .1.1.. =   20        w= 1111.. =   60
         26 :   a= ....11 =    3        w= 111..1 =   57
         27 :   a= ...11. =    6        w= 11..1. =   50
         28 :   a= ..11.. =   12        w= 1..1.. =   36
         29 :   a= .11... =   24        w= ..1..., =    8
         30 :   a= ..1.11 =   11        w= .1...1 =   17
         31 :   a= .1.11. =   22        w= 1...1. =   34
         32 :   a= ...111 =    7        w= ...1.1 =    5
         33 :   a= ..111. =   14        w= ..1.1. =   10
         34 :   a= .111.. =   28        w= .1.1.. =   20
         35 :   a= .1..11 =   19        w= 1.1..1 =   41

         36 :   a= .....1 =    1        w= .1..11 =   19 <-- period restarts
         37 :   a= ....1. =    2        w= 1..11. =   38
         38 :   a= ...1.. =    4        w= ..11.. =   12
```

Figure 39.6-A: Successive states of a FCSR with modulus $c = 37$.

```
1   class fcsr
2   {
3   public:
4       ulong a_;     // internal state (a_0*2**k modulo c),  1 <= a < c
5       ulong w_;     // word of the SRS,  1 <= w <= mask
6       ulong c_;     // a prime with primitive root 2, e.g. 37 = 1..1.1
7       ulong mask_;  // mask  e.g. (with above)    mask == 63 == 111111
8
9   public:
10      fcsr(ulong c)
11      {
12          c_ = c;
13          const ulong h = highest_one(c_);
14          mask_ = ( h | (h-1) );
15          set_a(1);
16      }
17
18      ~fcsr()  { ; }
19
20      ulong next()
21      {
22          a_ <<= 1;                   // a *= 2
23          if ( a_ > c_ )  a_ -= c_;   // reduce mod c
24
25          // update w:
26          w_ <<= 1;
27          w_ |= (a_ & 1);
28          w_ &= mask_;
29
30          return w_;
31      }
32
33    [--snip--]
34
```

```
35      void set_a(ulong a)
36      {
37          w_ = 0;
38          ulong t = c_;
39          while ( (t>>=1) )
40          {
41              if ( 0==(a & 1) )  a >>= 1;
42              else
43              {
44                  a = (a & c_) + ((a ^ c_) >> 1);
45              }
46          }
47          a_ = a;
48          next_w();
49      }
50
51      ulong get_a()  const  { return a_; }
52      ulong get_w()  const  { return w_; }
53  };
```

The routine is corresponds to the so-called Galois setup, described (for the LFSR) in section 39.2 on page 860, see also [131]. Figure 39.6-A shows the successive states of a FCSR with modulus $c = 37$. This is the output of [FXT: gf2n/fcsr-demo.cc]. Note that w does not run through all values smaller than c but through a subset of $c - 1$ distinct values smaller than 2^6.

```
    x:  p prime with 2 a primitive root, 2**x < p < 2**(x+1)
    1:  3
    2:  5
    3:  11 13
    4:  19 29
    5:  37 53 59 61
    6:  67 83 101 107
    7:  131 139 149 163 173 179 181 197 211 227
    8:  269 293 317 347 349 373 379 389 419 421 443 461 467 491 509
    9:  523 541 547 557 563 587 613 619 653 659 661 677 701 709 757
        773 787 797 821 827 829 853 859 877 883 907 941 947 1019
   10:  1061 1091 1109 1117 1123 1171 1187 1213 1229 1237 1259 1277
        1283 1291 1301 1307 1373 1381 1427 1451 1453 1483 1493 1499
        1523 1531 1549 1571 1619 1621 1637 1667 1669 1693 1733 1741
        1747 1787 1861 1867 1877 1901 1907 1931 1949 1973 1979 1987
        1997 2027 2029
```

Figure 39.6-B: List of primes $p < 2048$ where 2 is s primitive root.

A list of all primes less than 2048 for which 2 is a primitive root is shown in figure 39.6-B. The shown sequence is entry A001122 of [240]. For further information on the correspondence between LFSR and FCSR see [165].

39.7 Linear hybrid cellular automata (LHCA)

Linear hybrid cellular automata (LHCA) are 1-dimensional cellular automata (with 0 and 1 the only possible states) where two different rules are applied dependent of the position (therefore the 'hybrid' in the name). An implementation in C is [FXT: `lhca_next()` in bpol/lhca.h]:

```
1   inline ulong lhca_next(ulong x, ulong r, ulong m)
2   // LHCA := (1-dim) Linear Hybrid Cellular Automaton.
3   // Return next state (after x) of the LHCA with
4   // rule (defined by) r:
5   //   Rule 150 is applied for cells where r is one, rule 90 else.
6   //   Rule 150 :=  next(x) = x + leftbit(x) + rightbit(x)
7   //   Rule 90  :=  next(x) = leftbit(x) + rightbit(x)
8   // Length defined by m:
9   //   m has to be a burst of the n lowest bits (n: length of automaton)
10  {
11      r &= x;
12      ulong t = (x>>1) ^ (x<<1);
13      t ^= r;
14      t &= m;
15      return  t;
16  }
```

Note that the routine is branch free. Implementation in hardware is trivial as only neighboring bits interact.

The naming convention for the rules is as follows: draw a table of the eight possible states of a cell together with its neighbors then draw the new states below:

```
XXX  XX0  X0X  X00  0XX  0X0  00X  000
 0    X    0    X    X    0    X    0
```

Now read the lower row as a binary number, the result equals $01011010_2 = 90$, so this is rule 90. Rule 150 corresponds to $10010110_2 = 150$:

```
XXX  XX0  X0X  X00  0XX  0X0  00X  000
 X    0    0    X    0    X    X    0
```

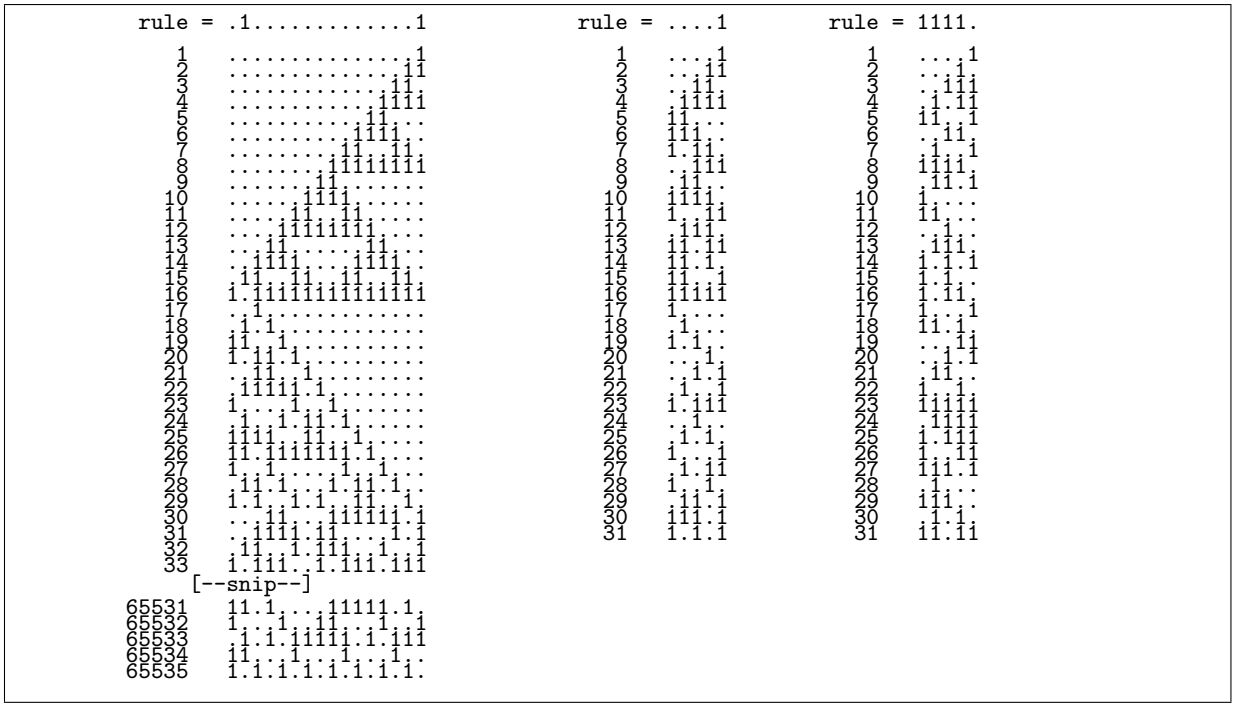

Figure 39.7-A: Partial run of a 16-bit LHCA (left) and complete runs of two 5-bit LHCAs (middle and right). These LHCAs have maximal periods.

A run of successive values for the length-16 weight-2 rule vector $r = 4001_{16}$ starting with 1 is shown on the left side of figure 39.7-A. For certain rule vectors r all $m = 2^n - 1$ nonzero values occur, the period is maximal. This is demonstrated in [FXT: gf2n/lhca-demo.cc], which for $n = 5$ and rule $r = 1$ gives the output shown in the middle of figure 39.7-A. Rule vectors with minimal weight that lead to maximal period are given in [78]. The list [FXT: `minweight_lhca_rule()` in bpol/lhcarule-minweight.cc] is taken from that source:

```
#define R1(n,s1)     (1UL<<s1)
#define R2(n,s1,s2)  (1UL<<s1) | (1UL<<s2)
extern const ulong minweight_lhca_rule[]=
// LHCA rules of minimum weight that lead to maximal period.
{
      0,  // (empty)
 R1(  1,  0),
 R1(  2,  0),
 R1(  3,  0),
 R2(  4,  0, 2),
 R1(  5,  0),
 R1(  6,  0),
 R1(  7,  2),
 R2(  8,  1, 2),
 R1(  9,  0),
 R2( 10,  1, 6),
 R1( 11,  0),
 R2( 12,  2, 6),
```

```
19    R1( 13,  4),
20    [--snip--]
21   };
```

Up to $n = 500$ there is always a rule with weight at most 2 that leads to the maximal period. The full list of these rules is given in [FXT: data/minweight-lhca-rules.txt].

39.7.1 Conversion to binary polynomials

To convert a length-n LHCA to a binary polynomial proceed as follows: initialize $p_{-1} := 0$, $p_0 := 1$ and iterate for $k = 1, 2, \ldots, n$:

$$p_k \quad := \quad (x + r_{k-1})\, p_{k-1} + p_{k-2} \tag{39.7-1}$$

where r_i denotes bit i of rule r. The degree of the returned polynomial p_n is n. An implementation of the algorithm is [FXT: `lhca2poly()` in bpol/lhca.h]:

```
1   inline ulong lhca2poly(ulong r, ulong n)
2   // Return binary polynomial p that corresponds
3   //  to the length-n LHCA rule r.
4   // p has degree n.
5   {
6       ulong p2 = 0,  p1 = 1;
7       while ( n-- )
8       {
9           ulong m = r & 1;
10          r >>= 1;
11          ulong p = (p1<<1) ^ p2;
12          if ( m )  p ^= p1;
13          p2 = p1;  p1 = p;
14
15      }
16      return  p1;
17  }
```

```
                        LHCA rule                          polynomial
  1:    r=..............................1  c=..............................11  r= [0]
  2:    r=..............................1  c=............................111  r= [0]
  3:    r=..............................1  c=..........................11.1  r= [0]
  4:    r=............................1.1  c=........................1..11  r= [0, 2]
  5:    r=..............................1  c=......................11.111  r= [0]
  6:    r=..............................1  c=.....................111..11  r= [0]
  7:    r=.............................1..  c=....................11.....1  r= [2]
  8:    r=............................11.  c=...................1...111.1  r= [1, 2]
  9:    r=..............................1  c=...................11.111..11  r= [0]
 10:    r=........................1....1.  c=.................1...11.1111  r= [1, 6]
 11:    r=..............................1  c=................11.1...111.1  r= [0]
 12:    r=..........................1...1.  c=................1.....11111.1  r= [2, 6]
 13:    r=........................1....  c=...............11..11.....111  r= [4]
 14:    r=..............................1  c=..............111..11......11  r= [0]
 15:    r=.............................1..  c=.............11.....1...1...1  r= [2]
 16:    r=...............1..............1  c=.............1..11.111.....111  r= [0, 14]
 17:    r=........................1....  c=............11..11..11.1....11  r= [4]
 18:    r=..............1..............1  c=............1..1..1.1.1.....1.1  r= [0, 16]
 19:    r=.............................1..  c=...........11.....111.1....11.1  r= [2]
 20:    r=............................11.  c=..........1...111.....1...1...1  r= [1, 2]
 21:    r=........................1.....1  c=.........1.11..1....1..1.1.1111  r= [0, 9]
 22:    r=........................1....  c=........1111..11...1.11.1111.11  r= [4]
 23:    r=..............................1  c........11.1...1.......111.1...1  r= [0]
 24:    r=...................1...1.......  c.......1.....1.1...1.1..11111..1  r= [7, 11]
 25:    r=........................1.......  c......11..11..11.........111..11  r= [8]
 26:    r=..............................1  c.....111..11.111..........11.111  r= [0]
 27:    r=............1.................1  c....1.111..1..111....11.1..111.1  r= [0, 19]
 28:    r=.............................1..  c...111111.1.1111........1.1.1111  r= [2]
 29:    r=..............................1  c=..11.111.....111.............111  r= [0]
 30:    r=..............................1  c=.111..11......11..............11  r= [0]
 31:    r=......................1.........  c=11.....................1...1...1  r= [10]
```

Figure 39.7-B: Minimum-weight LHCA rules and the corresponding binary polynomials.

The first minimum-weight LHCA rules and their binary polynomials are shown in figure 39.7-B. The table was created by the program [FXT: gf2n/lhca2poly-demo.cc].

For rules of maximal period the polynomials are primitive. In fact there is an isomorphism between the cycle structure of the successive states of an LHCA and its corresponding shift register.

An LHCA rule and its reverse give the identical polynomial.

The polynomials corresponding to a LHCA rule and its complement are either both reducible or irreducible: if $p(x)$ corresponds to the LHCA rule r then $p(x+1)$ corresponds to the rule that is the complement of r.

```
                    LHCA rule                          polynomial
 1:    r=.............................1  c=.............................11  r= [0]
 2:    r=.............................1  c=............................111  r= [0]
 3:    r=.............................1  c=...........................11.1  r= [0]
 4:    r=............................1.1  c=...........................1..11  r= [0, 2]
 5:    r=.............................1  c=..........................11.111  r= [0]
 6:    r=.............................1  c=.........................111..11  r= [0]
 7:    r=............................1..  c=.......................11.....1  r= [2]
 8:    r=...........................11.  c......................1...111.1  r= [1, 2]
 9:    r=.............................1  c.....................11.111..11  r= [0]
10:    r=...........................1111  c.....................1.111111.11  r= [0, 1, 2, 3]
11:    r=.............................1  c.....................11.1...111.1  r= [0]
12:    r=.........................1.11.  c..................11..1.1.....1  r= [1, 2, 4]
13:    r=...........................1..1  c.................1.11..1.1111.1  r= [0, 3]
14:    r=.............................1  c................111..11......11  r= [0]
15:    r=............................1..  c................11.....1...1...1  r= [2]
16:    r=.........................1.1.1  c..............11.1..1111..1.111  r= [0, 2, 4]
17:    r=............................11  c.............1.1111..11......11  r= [0, 1]
18:    r=.........................1.11.  c............11...11.1.1..1..1.1  r= [1, 2, 4]
19:    r=............................1..  c...........11.....111.1....11.1  r= [2]
20:    r=............................11.  c..........1...111.....1...1...1  r= [1, 2]
21:    r=..........................11..1.  c.........1111.11.1111.1..1.11.1  r= [1, 4, 5]
22:    r=.........................1.11  c.........11.1.111..11.11.11.1.11  r= [0, 1, 3]
23:    r=.............................1  c........11.1...1.......111.1...1  r= [0]
24:    r=.......................1.111.  c.......1.1..1.1....11....1..1..1  r= [1, 2, 3, 5]
25:    r=............................1.11  c......111....1.1...1.11.1....1.1  r= [0, 1, 3]
26:    r=.............................1  c.....111..11.111..........11.111  r= [0]
27:    r=...........................111..  c....1111..1.1111...1......1.1111  r= [2, 3, 4]
28:    r=............................1..  c...111111.1.1111........1.1.1111  r= [2]
29:    r=.............................1  c=..11.111.....111.............111  r= [0]
30:    r=.............................1  c=.111..11......11.............11  r= [0]
31:    r=...........................11.1.  c=11111.1.....1.11.......1....1.11  r= [1, 3, 4]
```

Figure 39.7-C: Low-bit LHCA rules and the corresponding binary polynomials.

Figure 39.7-C shows a list of LHCA rules where the highest bit lies in the lowest possible position. The list can be produced with [FXT: gf2n/lowbit-lhca-demo.cc]. For software implementation of LHCA that need more than one machine word low-bit rules are advantageous. A table of low-bit LHCA rules corresponding to primitive polynomials up to $n = 400$ is given in [FXT: data/lowbit-lhca-rules.txt]. The maximal value of a rule that occurs is $r = 1293$ for $n = 380$ so the table can be stored compactly.

Figure 39.7-D shows a list of LHCA rules that have minimal weight and the highest bit in the lowest possible position. The list can be produced with [FXT: gf2n/minweight-lowbit-lhca-demo.cc].

39.7.2 Conversion of irreducible binary polynomials to LHCAs

The LHCA corresponding to an irreducible binary polynomial $p(x)$ proceeds in two steps. Firstly, the quadratic equation (over GF(2^n))

$$Z^2 + (x^2 + x)\, p'(x)\, Z + 1 \quad = \quad 0 \mod p(x) \tag{39.7-2}$$

must be solved for Z. The algorithm is given in section 40.4 on page 886, the implementation is given in [FXT: bpol/bitpolmod-solvequadratic.cc]. The second step is a GCD computation. Set $Z_{n-1} := p$ and

```
                        LHCA rule                          polynomial
 1:    r=..............................1  c=..............................11  r= [0]
 2:    r=..............................1  c=.............................111  r= [0]
 3:    r=..............................1  c............................11.1  r= [0]
 4:    r............................1.1  c...........................1..11  r= [0, 2]
 5:    r..............................1  c.........................11.111  r= [0]
 6:    r..............................1  c.......................111..11  r= [0]
 7:    r............................1..  c.....................11.....1  r= [2]
 8:    r............................11.  c....................1...111.1  r= [1, 2]
 9:    r..............................1  c...................11.111..11  r= [0]
10:    r..........................1....1.  c.................1...11.1111  r= [1, 6]
11:    r..............................1  c.................11.1...111.1  r= [0]
12:    r........................1...1..  c................1.....11111.1  r= [2, 6]
13:    r............................1...  c................11..11.....111  r= [4]
14:    r..............................1  c..............111..11......11  r= [0]
15:    r............................1..  c.............11.....1...1...1  r= [2]
16:    r........................1.1....  c.............1..........111111  r= [4, 6]
17:    r............................1...  c............11..11..11.1....11  r= [4]
18:    r........................1...1.  c............1...11.11.1111111.1  r= [1, 5]
19:    r..............................1  c..........11.....111.1....11.1  r= [2]
20:    r............................11.  c..........1...111.....1...1...1  r= [1, 2]
21:    r........................11.....  c.........1.1...1......1.1.1...1  r= [5, 6]
22:    r............................1...  c.........1111..11...1.11.1111.11  r= [4]
23:    r..............................1  c........11.1...1.......111.1...1  r= [0]
24:    r..................1...1.......  c.......1.....1.1...1.1..11111..1  r= [7, 11]
25:    r.....................1........  c......11..11..11.........111..11  r= [8]
26:    r..............................1  c.....111..11.111..........11.111  r= [0]
27:    r........................1..1...  c....1.1......1..1..11111.....1.1  r= [3, 6]
28:    r..............................1  c...111111.1.1111........1.1.1111  r= [2]
29:    r..............................1  c..11.111.....111.............111  r= [0]
30:    r..............................1  c.111..11......11..............11  r= [0]
31:    r......................1........  c=11.....................1...1...1  r= [10]
```

Figure 39.7-D: Minimum-weight low-bit LHCA rules and the corresponding binary polynomials.

$Z_{n-2} := Z$, successively compute Z_{n-3}, Z_{n-4}, ..., Z_0 so that

$$\begin{aligned}
Z_{n-1} &= (x + r_0)\, Z_{n-2} + Z_{n-3} && \text{(39.7-3a)}\\
Z_{n-2} &= (x + r_1)\, Z_{n-3} + Z_{n-4} && \text{(39.7-3b)}\\
&\vdots\\
Z_2 &= (x + r_{n-3})\, Z_1 + 1 && \text{(39.7-3c)}\\
Z_1 &= (x + r_{n-2})\, Z_0 + 1 && \text{(39.7-3d)}\\
Z_0 &= (x + r_{n-1})\, 1 + 0 && \text{(39.7-3e)}
\end{aligned}$$

Each step consists of a computation of polynomial quotient and remainder (see section 38.1 on page 815). The vector $[r_{n-1}, r_{n-2}, \ldots, r_0]$ then gives the LHCA rule. The implementation is given [FXT: bpol/bitpol2lhca.cc]:

```
ulong
poly2lhca(ulong p)
// Return LHCA rule corresponding to the binary polynomial P.
// Must have: P irreducible.
{
    ulong dp = bitpol_deriv(p);
    const ulong h = bitpol_h(p);

    ulong b = dp;
    b ^= bitpolmod_times_x(b, p, h); // p' * (x+1)
    b = bitpolmod_times_x(b, p, h);  // p' * (x^2+x)

    ulong r0, r1;  // solutions of  1 + (p'*(x*x+x))*z + z*z == 0 modulo p
    bool q = bitpolmod_solve_quadratic(1, b, 1, r0, r1, p);
    if ( 0==q )  return 0;

    // GCD steps:
    ulong r = 0;  // rule vector
    ulong x = p,  y = r0;  // same result with r1
    while ( y )
    {
```

```
22          ulong tq, tr;
23          bitpol_divrem(x, y, tq, tr);
24          r <<= 1;
25          r |= (tq & 1);
26          x = y;
27          y = tr;
28      }
29
30      return r;
31  }
```

The described algorithm is given in the very readable paper [77] which is recommended for further studies. Note that the paper uses the reversed rule which can be obtained by inserting the line

```
r = revbin(r, bitpol_deg(p) );
```

just before the final return statement. The program [FXT: gf2n/poly2lhca-demo.cc] converts a given polynomial into the corresponding LHCA rule.

39.8 Additive linear hybrid cellular automata

39.8.1 Conversion into binary polynomials

The algorithm for the conversion of LHCA rules to binary polynomials is a special case of a general method for additive cellular automata. An automaton is called additive if, for all words a and b,

$$N(a) + N(b) \quad = \quad N(a + b) \tag{39.8-1}$$

where $N(x)$ is the next state after the state x and addition is bit-wise (XOR).

For additive automata the action of N on a binary word can be described by a matrix over GF(2): Let e_k be the word with bit number k the only set bit (e_k, $k = 0, \ldots, n-1$ is the canonical basis). The matrix whose k-th row is $N(e_k)$ is the matrix we seek. For example, for the LHCA with cyclic boundary condition (CLHCA) whose action N can be implemented as [FXT: bpol/clhca.h]

```
1  ulong clhca_next(ulong x, ulong r, ulong n)
2  {
3      r &= x;
4      ulong t = x ^ bit_rotate_right(x, 1, n);
5      t ^= r;
6      return  t;
7  }
```

we obtain for the rule $r := [r_0, r_1, \ldots, r_{n-1}]$ the matrix

$$M_r \quad := \quad \begin{bmatrix} s_0 & & & & & & & 1 \\ 1 & s_1 & & & & & & \\ & 1 & s_2 & & & & & \\ & & 1 & s_3 & & & & \\ & & & \ddots & \ddots & & & \\ & & & & 1 & s_{n-3} & & \\ & & & & & 1 & s_{n-2} & \\ & & & & & & 1 & s_{n-1} \end{bmatrix} \tag{39.8-2}$$

where s_k is the complement of r_k and blank entries denote zero. The binary polynomial corresponding to the automaton is the characteristic polynomial of the matrix M_r. We use the routine `bitmat_charpoly()` given in [FXT: bmat/bitmat-charpoly.cc] (which uses a reduction to the Hessenberg form as given in [90]):

```
1  inline ulong clhca2poly(ulong r, ulong n)
2  // Compute the binary polynomial corresponding
3  //   to the length-n CLHCA with rule r.
4  {
5      ALLOCA(ulong, M, n);
```

```
6       for (ulong k=0; k<n; ++k)
7       {
8           M[k] = clhca_next( 1UL<<k, r, n );
9       }
10      ulong c = bitmat_charpoly(M, n);
11      return c;
12  }
```

The routine computes the polynomial for any additive automaton if the call to `clhca_next()` is replaced by the update for the automaton. Note that with our particular example the matrix is already in Hessenberg form so we could directly call `bitmat_hessenberg2charpoly()`.

39.8.2 Properties of the CLHCA

```
          r = .111    r = 1.11    r = 11.1    r = 111.

      1:      ...1        ...1        ...1        ...1
      2:      1...        1...        1...        1..1
      3:      11..        .1..        .1..        11.1
      4:      111.        .11.        ..1.        1111
      5:      1111        .111        ..11        111.
      6:      .111        1111        1.11        .111
      7:      1.11        1.11        1111        1.1.
      8:      .1.1        11.1        11.1        .1.1
      9:      1.1.        1.1.        111.        1.11
     10:      11.1        .1.1        .1.1        11..
     11:      .11.        111.        1.1.        .11.
     12:      ..11        ..11        .111        ..11
     13:      1..1        1..1        1..1        1...
     14:      .1..        11..        11..        .1..
     15:      ..1.        ..1.        .11.        ..1.
```

Figure 39.8-A: Rules of identical weight lead to essentially identical CLHCA: all tracks in the successive states of all weight-3 (length-4) automata are cyclic shifts of each other.

```
     k      r =  CLHCA rule            c = binary polynomial
     0:     r = ................       c = 11..............1.
     1:     r = ...............1       c = 1...............11
     2:     r = ..............11       c = 1111111111111111.1
     3:     r = .............111       c = 1.1.1.1.1.1.1.1..1  P
     4:     r = ............1111       c = 11..11..11..11...1
     5:     r = ...........11111       c = 1...1...1...1....1  P
     6:     r = ..........111111       c = 1111....1111.....1  P
     7:     r = .........1111111       c = 1.1.....1.1......1
     8:     r = ........11111111       c = 11......11.......1
     9:     r = .......111111111       c = 1.......1........1
    10:     r = ......1111111111       c = 11111111.........1
    11:     r = .....11111111111       c = 1.1.1.1..........1  P
    12:     r = ....111111111111       c = 11..11...........1  P
    13:     r = ...1111111111111       c = 1...1............1
    14:     r = ..11111111111111       c = 1111.............1  P
    15:     r = .111111111111111       c = 1.1..............1
    16:     r = 1111111111111111       c = 11...............1
    17:     r = .111111111111111       c = 1................1
```

Figure 39.8-B: Binary polynomials (right) corresponding to rule vectors for the length-17 cyclic linear hybrid cellular automata (left) with k bits set. Automata with maximal period are marked with '`P`'.

The polynomials for the given automata depend only on the number of bits set in the rule r. Figure 39.8-A shows the successive states for all length-4 CLHCA with rules of weight 3. As there are essentially only $n+1$ different automata of length n we only need to investigate the rules with the lowest k bits set for $k = 0, 1, \ldots, n$. The polynomials for the length-17 automata are shown in figure 39.8-B. The polynomials can be computed with the simple routine

```
1   inline ulong clhca2poly(ulong r, ulong n)
2   {
3       ulong c = 1UL << n;
4       for (ulong k=0; k<n; ++k)  if ( 0==(r & (1UL<<k)) )  c ^= (c>>1);
5       c ^= 1;
6       return c;
7   }
```

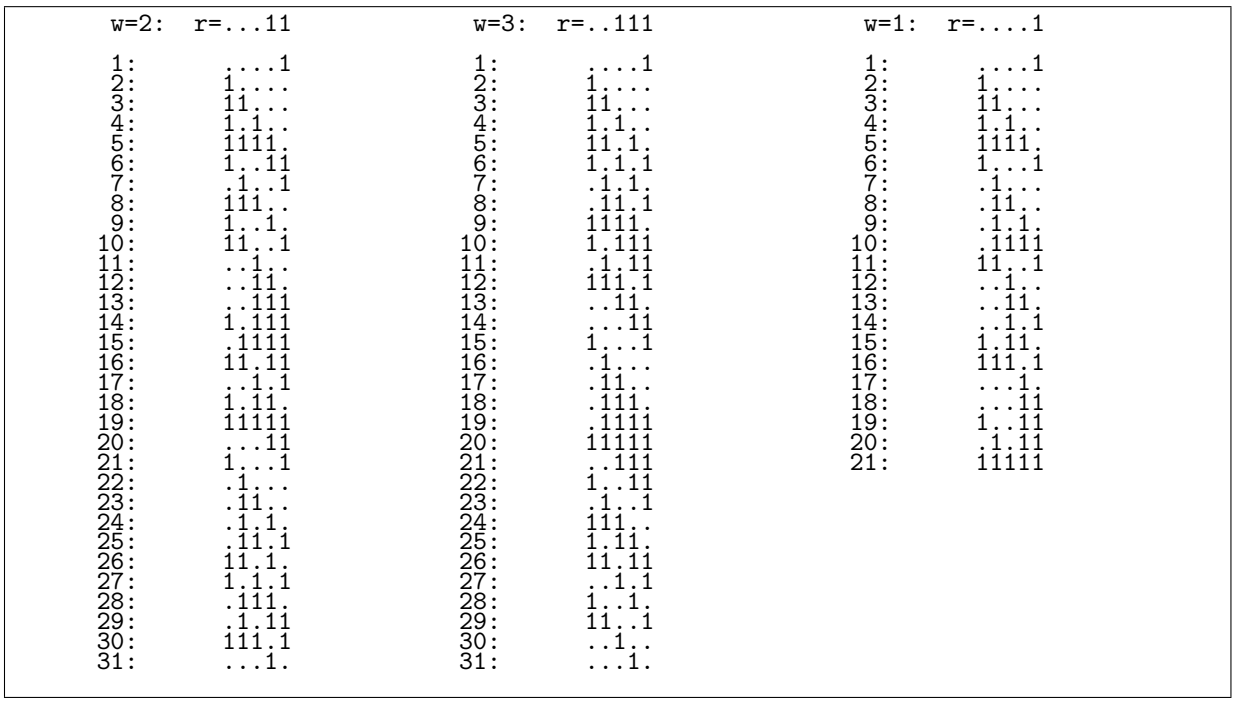

```
w=2:  r=...11          w=3:  r=..111          w=1:  r=....1

 1:      ....1          1:      ....1          1:      ....1
 2:      1....          2:      1....          2:      1....
 3:      11...          3:      11...          3:      11...
 4:      1.1..          4:      1.1..          4:      1.1..
 5:      1111.          5:      11.1.          5:      1111.
 6:      1..11          6:      1.1.1          6:      1...1
 7:      .1..1          7:      .1.1.          7:      .1...
 8:      111..          8:      .11.1          8:      .11..
 9:      1..1.          9:      1111.          9:      .1.1.
10:      11..1         10:      1.111         10:      .1111
11:      ..1..         11:      .1.11         11:      11..1
12:      ..11.         12:      111.1         12:      ..1..
13:      ..111         13:      ..11.         13:      ..11.
14:      1.111         14:      ...11         14:      ..1.1
15:      .1111         15:      1...1         15:      1.11.
16:      11.11         16:      .1...         16:      111.1
17:      ..1.1         17:      .11..         17:      ...1.
18:      1.11.         18:      .111.         18:      ...11
19:      11111         19:      .1111         19:      1..11
20:      ...11         20:      11111         20:      .1.11
21:      1...1         21:      ..111         21:      11111
22:      .1...         22:      1..11
23:      .11..         23:      .1..1
24:      .1.1.         24:      111..
25:      .11.1         25:      1.11.
26:      11.1.         26:      11.11
27:      1.1.1         27:      ..1.1
28:      .111.         28:      1..1.
29:      .1.11         29:      11..1
30:      111.1         30:      ..1..
31:      ...1.         31:      ...1.
```

Figure 39.8-C: The length-5, weight-w CLHCA has maximal period for $w = 2$ and $w = 3$. The successive states are shown for both automata (left and middle). The weight-1 rule leads to a period of 21 (right).

```
 n  w:     c = polynomial

 2  1:     c = 111
 3  1:     c = 1.11
 4  3:     c = 11..1
 5  2:     c = 1111.1
 6  1:     c = 11..111
 7  1:     c = 1.1.1.11
[8]
 9  4:     c = 11..11...1
10  3:     c = 11111111..1
11  2:     c = 11......11.1
[12,13,14]
15  4:     c = 1111....1111...1
[16]
17  3:     c = 1.1.1.1.1.1.1.1..1
18  7:     c = 1111....1111......1
[19]
20  3:     c = 11............11..1
21  2:     c = 1111..........1111.1
22 21:     c = 11..................1
23  5:     c = 1.1............1.1....1
[24]
25  3:     c = 1.1.1.1.........1.1.1.1..1
[26,27]
28  9:     c = 1111............1111.........1
29  2:     c = 1111....1111....1111....1111.1
[30]
```

Figure 39.8-D: The polynomials (right) corresponding rules of lowest weight (w) so that the length-n ($n \leq 30$) CLHCA has maximal period.

If we write $C_{n,k}(x)$ for the polynomial for the length-n, k-bit rule then $C_{n,n-k}(x) = C_{n,k}(x+1)$. Indeed, they can be given in the closed form $C_{n,k}(x) = 1 + x^k\,(1+x)^{n-k}$.

With $n = 5$ there are just two rules that lead to maximal periods, $r = [0, 0, 0, 1, 1]$ (weight 2), and $r = [0, 0, 1, 1, 1]$ (weight 3). The successive states for both rules are shown in figure 39.8-C. The polynomials corresponding to the rules of minimal weight for all length-n automata where $n \leq 30$ are given in figure 39.8-D. The sequence of values n so that a primitive length-n CLHCA exists starts:

```
2, 3, 4, 5, 6, 7, 9, 10, 11, 15, 17, 18, 20, 21, 22, 23, 25, 28, 29, 31, 33,
35, 36, 39, 41, 47, 49, 52, 55, 57, 58, 60, 63, ...
```

It coincides with entry A073726 of [240], values n such that there is a primitive trinomial of degree n. The sequence was computed with the program [FXT: gf2n/clhca-demo.cc]. A list of CLHCA rules with maximal period is given in [FXT: data/clhca-rules.txt].

Chapter 40

Binary finite fields: $\mathrm{GF}(2^n)$

This chapter introduces the binary finite fields $\mathrm{GF}(2^n)$. The polynomial representation is used for stating some basic properties. An introduction of the representation by normal bases follows. Certain normal bases are advantageous for hardware implementations of the arithmetic algorithms. Finally, several ways of computing the number of irreducible binary normal polynomials are given.

Binary finite fields are important for applications like error correcting codes and cryptography.

The underlying arithmetical algorithms are given in chapter 38.

40.1 Arithmetic and basic properties

In chapter 25 we have met the finite fields $\mathbb{Z}/p\mathbb{Z} = \mathrm{GF}(p)$ for p a prime. The 'GF' stands for *Galois Field*, another symbol often used is $\mathbb{F}_p$. The arithmetic in $\mathrm{GF}(p)$ is the simply the arithmetic modulo p.

There are more finite fields: for every prime p there are fields with $Q = p^n$ elements for all $n \geq 1$. All elements in a finite field $\mathrm{GF}(p^n)$ can be represented as polynomials modulo an degree-n irreducible polynomial C with coefficients over the field $\mathrm{GF}(p)$. The arithmetic to be used is polynomial arithmetic modulo C. As in general there is more than one irreducible polynomial of degree n it might seem that there is more than one field $\mathrm{GF}(Q)$ for given $Q = p^n$. There isn't. Using different polynomials as modulus leads to isomorphic representations of the same field. The field $\mathrm{GF}(p^n)$ is called an *extension field* of $\mathrm{GF}(p)$. The field $\mathrm{GF}(p)$ is called the *ground field* of $\mathrm{GF}(p^n)$.

When speaking about an element of $\mathrm{GF}(Q)$ one can think of a polynomial modulo some fixed irreducible polynomial C (modulus). For example, the product of two elements can be computed as the polynomial product modulo C. For the equivalent construction using the polynomial $x^2 + 1$ with real coefficients that leads to the complex numbers, see section 37.12 on page 793.

The elements zero (the neutral element of addition) and one (neutral element of multiplication) are the constant polynomials with constant zero and one, respectively. This does not depend on the choice of the modulus.

Multiplication with an element of the ground field is called *scalar multiplication*. In this section an element of the ground field is denoted by u. Scalar multiplication corresponds to the multiplication of every coefficient of (the polynomial representing) the field element a by u.

We restrict our attention mostly to $Q = 2^n$, that is, the *binary finite fields* $\mathrm{GF}(2^n)$ in what follows as we have seen the algorithms for the underlying arithmetic.

40.1.1 Characteristic and linear functions

If we add any element of the field GF(p^n) p times to zero the result will be zero. One calls p the *characteristic* of the field. For infinite fields such as $\mathbb{C}$ the characteristic is said to be zero (while it should really be ∞). Note that the notion "multiplication is repeated addition" is meaningless in extension fields.

For GF(p^n) one has

$$(u+v)^p \quad = \quad u^p + v^p \tag{40.1-1}$$

because the binomial coefficients $\binom{p}{k}$ are divisible by p for $k = 1, 2, \ldots, p-1$. For GF(2^n):

$$(u+v)^2 \quad = \quad u^2 + v^2 \tag{40.1-2}$$

We call a function f *linear* if the relation

$$f(u_1 \cdot a + u_2 \cdot b) \quad = \quad u_1 \cdot f(a) + u_2 \cdot f(b) \tag{40.1-3}$$

holds for u_1 and u_2 from the ground field. The linear functions in GF(p^n) are of the form

$$f(x) \quad = \quad \sum_{k=0}^{n-1} u_k \cdot x^{p^k} \tag{40.1-4}$$

where the u_k are again in the ground field. Linear functions can be computed using lookup tables. In GF(2^n) these are all functions of the form

$$f(x) \quad = \quad \sum_{k=0}^{n-1} u_k \cdot x^{2^k} \tag{40.1-5}$$

40.1.2 Squaring

Squaring (and raising to any power 2^k) is a linear operation in GF(2^n). The linearity can be used to accelerate the computation of squares. Write

$$\begin{aligned} \left(u_0 + u_1\, x + u_2\, x^2 + \ldots + u_{n-1}\, x^{n-1}\right)^2 \quad &= \quad u_0 + u_1\, x^2 + u_2\, x^4 + \ldots + u_{n-1}\, x^{2\,(n-1)} \\ &=: \quad u_0\, s_0 + u_1\, s_1 + u_2\, s_2 + \ldots u_{n-1}\, s_{n-1} \end{aligned} \tag{40.1-6}$$

One has to precompute and store the values $s_i = x^{2\,i} \bmod C$ for $i = 0, 1, 2, \ldots, n-1$. For successive square computations it is only necessary to add (that is, XOR) those s_i corresponding to nonzero u_i. For example, with $n = 13$ and the polynomial modulus $C = x^{13} + x^4 + x^3 + x^1 + 1$ one obtains the table

```
 S0 =  ............1
 S1 =  ...........1.
 S2 =  ..........1..
 S3 =  .........1...
 S4 =  ........1....
 S5 =  .......1.....
 S6 =  ......1......
 S7 =  .....11.11...
 S8 =  ....11.11....
 S9 =  ...11.11.....
S10 =  ..11.11......
S11 =  .1.11...11.11
S12 =  .11....1.11.1.
```

The squares $s_0, s_1 \ldots s_{\lfloor (n-1)/2 \rfloor}$ have the trivial form $s_i = x^{2\,i}$ which can be used to further accelerate the computation.

40.1.3 Computation of the trace

The *trace* of an element u in GF(2^n) is defined as

$$\mathrm{Tr}(a) \quad := \quad a + a^2 + a^4 + a^8 + \ldots + a^{2^{n-1}} \quad = \sum_{j=0}^{n-1} a^{2^j} \tag{40.1-7}$$

The trace of any element is either zero or one. The trace of zero is always zero. The trace of one in GF(2^n) equals one for odd n, else zero, that is, $\mathrm{Tr}(1) = n \bmod 2$. Exactly half of the elements have trace one. The trace of the sum of two elements is the sum of the traces of the elements:

$$\mathrm{Tr}(a+b) \quad = \quad \mathrm{Tr}(a) + \mathrm{Tr}(b) \tag{40.1-8}$$

The trace of the square of an element is the trace of that element:

$$\mathrm{Tr}(a^2) \quad = \quad \mathrm{Tr}(a) \tag{40.1-9}$$

With u zero or one (an element of the ground field GF(2)) one has

$$\mathrm{Tr}(u \cdot a) \quad = \quad u \cdot \mathrm{Tr}(a) \tag{40.1-10}$$

Thereby, for u_1 and u_1 from the ground field,

$$\mathrm{Tr}(u_1 \cdot a + u_2 \cdot b) \quad = \quad u_1 \cdot a + u_2 \cdot b \tag{40.1-11}$$

That is, the Tr function is linear.

A fast algorithm to compute the trace uses the *trace vector*, a precomputed table $t_i = \mathrm{Tr}(x^i)$ for $i = 0, 1, 2, \ldots, n-1$, and the linearity of the trace

$$\mathrm{Tr}(u) \quad = \quad \sum_{i=0}^{n-1} u_i \, t_i \tag{40.1-12}$$

where the u_i are zero or one. Thereby, when the polynomial fits into a single binary word, precompute the trace vector `tv` whose bits are the traces of the powers of x and later compute the trace of an element via [FXT: `gf2n_fast_trace()` in bpol/gf2n-trace.h]:

```
1  inline ulong gf2n_fast_trace(ulong a, ulong tv)
2  // Fast trace computation of the trace of a
3  //  using the precalculated table tv.
4  { return  parity( a & tv ); }  // scalar product over GF(2)
```

Given the trace vector it is also easy to find elements of trace zero or one by simply taking the lowest unset or set bit of the vector, respectively. There are polynomials such that the trace vector contains just one nonzero bit, see section 884.

40.1.4 Inverse and square root

The number of elements in GF(Q) equals Q. For any element $a \in \mathrm{GF}(Q)$ one has

$$a^Q \quad = \quad a \tag{40.1-13}$$

and thereby $a^{Q-1} = 1$. So we can compute the *inverse* a^{-1} of a nonzero element a as

$$a^{-1} \quad = \quad a^{Q-2} \tag{40.1-14}$$

We have see this technique of inversion by exponentiation is section 37.6.4 on page 768 which is the special case GF(p).

All elements except zero are invertible in a field. That is, the number of invertible elements (units) in GF(Q) equals $|\mathrm{GF}(Q)^*| = Q - 1 = p^n - 1$

Every element a of GF(2^n) has a unique *square root* s that can easily be computed as

$$s \;=\; a^{Q/2} \;=\; a^{2^{n-1}} \tag{40.1-15}$$

It can be computed by squaring the element $n-1$ times. But the square root is a linear function, so we can again apply table lookup methods.

A method that uses the precomputed value $\sqrt{x}$ is described in [118]: for an element $a = \sum_k a_k\, x^k$ we have

$$\sqrt{a} \;=\; \sum_{k\ \mathrm{even}} a_k\, x^{k/2} + \sqrt{x} \sum_{k\ \mathrm{odd}} a_k\, x^{(k-1)/2} \tag{40.1-16}$$

The only nontrivial operation is the multiplication with $\sqrt{x}$. The method is easily generalized for higher roots: for example, for the third root define $S(m) := \sum_{k\equiv m \bmod 3} a_k\, x^{(k-m)/3}$, then

$$\sqrt[3]{a} \;=\; S(0) + \sqrt[3]{x}\, S(1) + \left(\sqrt[3]{x}\right)^2 S(2) \tag{40.1-17a}$$

$$\;=\; S(0) + \sqrt[3]{x}\, \left[S(1) + \sqrt[3]{x}\, S(2)\right] \tag{40.1-17b}$$

40.1.5 Order and primitive roots

The *order* of an element a is the least positive exponent r so that $a^r = 1$. The maximal order of an element in GF(2^n) equals $2^n - 1 = Q - 1$. An element of maximal order is called a *generator* (or *primitive root*) as its powers 'generate' all nonzero elements in the field. The order of an given element a in GF(2^n) can be computed like

```
1   function order(a, n)
2   {
3       if a==0 then  return 0  // a not a unit
4
5       h := 2**n - 1  // number of units
6       e := h
7       {np, p[], k[]} := factorization(h)  // h==product(i=0..np-1, p[i]**k[i])
8
9       for i:=0 to np-1
10      {
11          f := p[i]**k[i]
12          e := e / f
13          g1 := a**e  // modulo polynomial
14          while g1!=1
15          {
16              g1 := g1**p[i]  // modulo polynomial
17              e := e * p[i]
18              p[i] := p[i] - 1
19          }
20      }
21
22      return e
23  }
```

The C++ implementation is [FXT: `gf2n_order()` in bpol/gf2n-order.cc]. We have seen a very similar algorithm in section 38.5.2 on page 835. Indeed, with just minimal changes we obtain a pari/gp implementation to compute the order of an element g in GF(2^n):

```
1   polorder(p, g='x) =
2   /* Order of g modulo p (p irreducible over GF(2)) */
3   {
4     local(g1, te, tp, tf, tx);
5     p *= Mod(1,2);
6     te = nn_;
7     for(i=1, np_,
8       tf = vf_[i];  tp = vp_[i];  tx = vx_[i];
9       te = te / tf;
```

```
10      g1 = Mod(g, p)^te;
11      while ( 1!=g1,
12        g1 = g1^tp;
13        te = te * tp;
14      );
15    );
16    return( te );
17  }
```

The only difference to the original algorithm is that g can have other values than (the polynomial) x. The function uses the following global variables that must be set up before call:

```
1   nn_ = 0;  /* max order  =  2^n-1 */
2   np_ = 0;  /* number of primes in factorization */
3   vp_ = []; /* vector of primes */
4   vf_ = []; /* vector of factors (prime powers) */
5   vx_ = []; /* vector of exponents */
```

To check whether g is a primitive root modulo p use the function

```
1   polmaxorder_q(p, g='x) =
2   /* Whether order of g modulo p is maximal  (p irreducible over GF(2)) */
3   /* Early-out variant */
4   {
5     local(g1, te, tp, tf, tx, ct);
6     p *= Mod(1,2);
7     te = nn_;
8     for(i=1, np_,
9       tf = vf_[i];  tp = vp_[i];  tx = vx_[i];
10      te = te / tf;
11      g1 = Mod(g, p)^te;
12      ct = 0;
13      while ( 1!=g1,
14        g1 = g1^tp;
15        te = te * tp;
16        ct = ct + 1;
17      );
18      if ( ct<tx,  return(0) );
19    );
20    return(1);
21  }
```

Let a be an element of order r, then the order of a^k is

$$\mathrm{ord}(a^k) \quad = \quad \frac{r}{\gcd(k,r)} \tag{40.1-18}$$

The order remains unchanged if $\gcd(k,r) = 1$. Let $N = 2^n - 1$, and g a generator of $GF(2^n)$, then all $\varphi(N)$ generators can obtained as g^k where $\gcd(k,N) = 1$. For N a (Mersenne-) prime the order of all invertible elements except 1 is N.

40.1.6 Implementation

A C++ class for computations in the fields $\mathrm{GF}(2^n)$ with n not greater than `BITS_PER_LONG` is [FXT: `class GF2n` in bpol/gf2n.h]:

```
1   class GF2n
2   // Implementation of binary finite fields GF(2**n)
3   //   with arithmetic operations on it.
4   {
5   public:
6       ulong v_;
```

The static (that is, class global) elements support the computations:

```
1   public:
2       static ulong n_;  // the 'n' in GF(2**n)
3       static ulong c_;  // polynomial modulus
4       static ulong h_;  // auxiliary bitmask for computations
5       static ulong mm_; // 2**n - 1 == max order (a Mersenne number)
```

```
6       static ulong g_;  // a generator (element of maximal order)
7       static ulong tv_; // trace vector
8       static ulong sqr_tab[BITS_PER_LONG]; // table for fast squaring
9       static factorization mfact_; // factorization of max order
10      static char* pc_; // chars to print zero and one: e.g.  "01" or ".1"
11      [--snip--]
12      static GF2n zero; // zero (neutral element wrt. addition) in GF(2**n)
13      static GF2n one;  // one (neutral element wrt. multiplication) in GF(2**n)
14      static GF2n tr1e; // an element with trace == 1
```

Note all data is public, making many methods '`get_something()`' unnecessary. You can also modify to the data. Expect funny results if you do. The constructors from other types are '`explicit`' to avoid surprises:

```
1   public:
2       explicit GF2n()  { ; }
3       explicit GF2n(const ulong i) : v_(i & mm_)  { ; }
4       GF2n(const GF2n &g) : v_(g.v_)  { ; }
5       ~GF2n()  { ; }
```

Before doing anything one has to call the initializing function [FXT: `GF2n::init()` in bpol/gf2n.cc]:

```
1   // if INIT_ASSERT is defined, asserts are C asserts,
2   // else init() returns false if one of the tests fail:
3   #define INIT_ASSERT
4
5   bool // static
6   GF2n::init(ulong n, ulong c/*=0*/, bool normalq/*=0*/, bool trustme/*=0*/)
7   // Initialize class GF(2**n) for 0<n<=BITS_PER_LONG.
8   // If an irreducible polynomial c is supplied it is used as modulus,
9   // else a primitive polynomial of degree n is used.
10  // Irreducibility of c is asserted for deg(c)<BITS_PER_LONG.
11  // When normalq is set a primitive normal polynomial is used,
12  // if in addition c is supplied, normality of c is asserted.
13  // When trustme is set the asserts are omitted.
14  {
15    [--snip--]
16          if ( n_ < BITS_PER_LONG )  // test only works for polynomials that fit into words
17          {
18              if ( ! trustme )
19              {
20  #ifdef INIT_ASSERT
21                  jjassert( bitpol_irreducible_q(c_, h_) );
22  #else
23                  if ( ! bitpol_irreducible_q(c_, h_) )  return false;
24  #endif
25              }
26          }
27    [--snip--]
28  }
```

The class defines all the standard operators like the binary operators '`+`' and '`-`' (which are the same operation in GF(2^n)), '`*`' and '`/`', the comparison operators '`==`' and '`!=`', also the computation of inverse, powering, order, and trace. The algorithms used for the arithmetic operations are described in section 38.3 on page 827. We give the method for the inverse and the arithmetic shortcut operators as examples:

```
1       GF2n inv()  const
2       {
3           GF2n z;
4           z.v_= bitpolmod_inverse(v_, GF2n::c_);
5           return z;
6       }
7       [--snip--]
8
9       friend inline GF2n & operator += (GF2n &z, const GF2n &f)
10      { z.v_ ^= f.v_;  return z; }
11
12      friend inline GF2n & operator -= (GF2n &z, const GF2n &f)
13      { z.v_ ^= f.v_;  return z; }
14
15      friend inline GF2n & operator *= (GF2n &z, const GF2n &f)
16      { z.v_ = bitpolmod_mult(z.v_, f.v_, GF2n::c_, GF2n::h_);  return z; }
17
18      friend inline GF2n & operator /= (GF2n &z, const GF2n &f)
19      { z *= f.inv();  return z; }
```

```
    n = 4  GF(2^n)
    c = 1..11 == x^4 + x + 1  (polynomial modulus)
    mm= .1111   == 15  =  3 * 5  (maximal order)
    h = .1... (aux. bitmask)
    g = ...1. (element of maximal order)
    tv= .1... (traces of x^i)
  tr1e= .1... (element with trace=1)

       k  :    f:=g**k  Tr(f)    ord(f)     f*f    sqrt(f)
     .... :      ...1     0        1       ...1     ...1
     ...1 :      ..1.     0       15       .1..     .1.1
     ..1. :      .1..     0       15       ..11     ..1.
     ..11 :      1...     1        5       11..     1.1.
     .1.. :      ..11     0       15       .1.1     .1..
     .1.1 :      .11.     0        3       .111     .111
     .11. :      11..     1        5       1111     1...
     .111 :      1.11     1       15       1..1     111.
     1... :      .1.1     0       15       ..1.     ..11
     1..1 :      1.1.     1        5       1...     1111
     1.1. :      .111     0        3       .11.     .11.
     1.11 :      111.     1       15       1.11     11.1
     11.. :      1111     1        5       1.1.     11..
     11.1 :      11.1     1       15       111.     1..1
     111. :      1..1     1       15       11.1     1.11
```

Figure 40.1-A: Powers of the generator $g = x$ in GF(2^4) when a primitive polynomial is used as a modulus.

```
    n = 4  GF(2^n)
    c = 11111 [normal] [NON-primitive]
      == x^4 + x^3 + x^2 + x + 1  (polynomial modulus)
    mm= .1111   == 15  =  3 * 5  (maximal order)
    h = .1... (aux. bitmask)
    g = ...11 (element of maximal order)
    tv= .111. (traces of x^i)
  tr1e= ...1. (element with trace=1)

       k  :    f:=g**k  Tr(f)    ord(f)     f*f    sqrt(f)
     .... :      ...1     0        1       ...1     ...1
     ...1 :      ..11     1       15       .1.1     1..1
     ..1. :      .1.1     1       15       111.     ..11
     ..11 :      1111     1        5       1...     .1..
     .1.. :      111.     1       15       1..1     .1.1
     .1.1 :      11.1     0        3       11..     11..
     .11. :      1...     1        5       ..1.     1111
     .111 :      .111     0       15       1.1.     1.11
     1... :      1..1     1       15       ..11     111.
     1..1 :      .1..     1        5       1111     ..1.
     1.1. :      11..     0        3       11.1     11.1
     1.11 :      1.11     0       15       .111     .11.
     11.. :      ..1.     1        5       .1..     1...
     11.1 :      .11.     0       15       1.11     1.1.
     111. :      1.1.     0       15       .11.     .111
```

Figure 40.1-B: Powers of a generator in GF(2^4) when a non-primitive polynomial is used as a modulus.

As a simple demonstration of the usage, the program [FXT: gf2n/gf2n-demo.cc] prints the successive powers of a primitive root, and their squares and square roots. By default computations in GF(2^4) are shown, both for a primitive polynomial modulus (figure 40.1-A), and for a non-primitive polynomial modulus (figure 40.1-B).

40.2 Minimal polynomials

The *minimal polynomial* of an element a in GF(2^n) is defined as the polynomial of least degree which has a as a root. The minimal polynomial can be computed as the product

$$p_a(x) \;:=\; \prod_{k=0}^{r-1} \left(x - a^{2^k}\right) \tag{40.2-1}$$

where r is the smallest positive integer so that $a^{2^r} = a$. The minimal polynomial of any element is irreducible and its degree is a divisor of n.

By construction the zeros of the polynomial are $a, a^2, a^4, a^8, \ldots, a^{2^{r-1}}$. Note that $\left(a^{2^{r-1}}\right)^2 = a$, that is, $p_a = p_{a^2} = p_{a^4} = \ldots = p_{a^{2^{r-1}}}$. The elements $a^2, a^4, \ldots$ are called the *conjugates* of a. For the field GF(p^n) the minimal polynomial has the form $\prod_{k=0}^{r-1} \left(x - a^{p^k}\right)$.

From the definition it can be seen that the coefficients of the minimal polynomial lie in GF(2^n), however, all of them are either zero or one. Thereby the computation has to be carried out using the arithmetic in GF(2^n) but the final result is a binary polynomial [FXT: bpol/gf2n-minpoly.cc]:

```
1   ulong
2   gf2n_minpoly(GF2n a, ulong &bp)
3   // Compute the minimal polynomial p(x) of a \in GF(2**n).
4   // Return the degree of p().
5   // The polynomial p() is written to bp
6   {
7       GF2n p[BITS_PER_LONG+1];
8       ulong n = GF2n::n_;
9       for (ulong k=0; k<=n; ++k)  p[k] = 0;
10      p[0] = 1;
11      ulong d;
12      GF2n s = a;
13      for (d=1; d<=n; ++d)
14      {
15          for (ulong k=d; 0!=k; --k)  p[k] = p[k-1];
16          p[0] = 0;
17          for (ulong k=0; k<d; ++k)  p[k] += p[k+1] * s;
18          s = s.sqr();
19          if ( s == a )  break;
20      }
21
22      // Here all coefficients are either zero or one,
23      //  so we can fill them into a binary word:
24      ulong p2 = 0;
25      for (ulong j=0; j<=d; ++j)  p2 |= (p[j].v_ << j);
26      bp = p2;
27
28      return  d;
29  }
```

A version of the routine that does not depend on the class `GF2n` is given in [FXT: bpol/bitpolmod-minpoly.h]. The program [FXT: gf2n/gf2n-minpoly-demo.cc] prints the minimal polynomials for the powers of a primitive element g, see figure 40.2-A. The polynomials (of maximal degree) that are non-primitive are marked by an '`N`'. The minimal polynomials for g^k are non-primitive or of degree smaller than n whenever $\gcd(k, 2^n - 1) \neq 1$.

Let C be a irreducible polynomial of degree n and a an element so that $r = \mathrm{ord}_C(a)$ is the order of a modulo C. Then the order of x modulo the minimal polynomial of a is also r. Thereby a primitive poly-

```
n = 6  GF(2^n)
c = 1....11 == x^6 + x + 1  (polynomial modulus)
mm= .111111   == 63  =  3^2 * 7  (maximal order)

 k          :    f:=g**k  ord(f) Tr(f)  p:=minpoly(f)  deg(p)
 0 = ......  :     .....1     1    0       ......11       1
 1 = .....1  :     ....1.    63    0       .1....11       6
 2 = ....1.  :     ...1..    63    0       .1....11       6
 3 = ....11  :     ..1...    21    0       .1.1.111       6 N
 4 = ...1..  :     .1....    63    0       .1....11       6
 5 = ...1.1  :     1.....    63    1       .11..111       6
 6 = ...11.  :     ....11    21    0       .1.1.111       6 N
 7 = ...111  :     ...11.     9    0       .1..1..1       6 N
 8 = ..1...  :     ..11..    63    0       .1....11       6
 9 = ..1..1  :     .11...     7    0       ....11.1       3
10 = ..1.1.  :     11....    63    1       .11..111       6
11 = ..1.11  :     1...11    63    1       .11.11.1       6
12 = ..11..  :     ...1.1    21    0       .1.1.111       6 N
13 = ..11.1  :     ..1.1.    63    0       .1.11.11       6
14 = ..111.  :     .1.1..     9    0       .1..1..1       6 N
15 = ..1111  :     1.1...    21    1       .111.1.1       6 N
16 = .1....  :     .1..11    63    0       .1....11       6
17 = .1...1  :     1..11.    63    1       .11..111       6
18 = .1..1.  :     ..1111     7    0       ....11.1       3
19 = .1..11  :     .1111.    63    0       .1.11.11       6
20 = .1.1..  :     1111..    63    1       .11..111       6
21 = .1.1.1  :     111.11     3    1       .....111       2
22 = .1.11.  :     11.1.1    63    1       .11.11.1       6
23 = .1.111  :     1.1..1    63    1       .111..11       6
24 = .11...  :     .1...1    21    0       .1.1.111       6 N
25 = .11..1  :     1...1.    63    1       .11.11.1       6
26 = .11.1.  :     ...111    63    0       .1.11.11       6
27 = .11.11  :     ..111.     7    0       ....1.11       3
28 = .111..  :     .111..     9    0       .1..1..1       6 N
29 = .111.1  :     111...    63    1       .111..11       6
30 = .1111.  :     11..11    21    1       .111.1.1       6 N
31 = .11111  :     1..1.1    63    1       .11....1       6
32 = 1.....  :     ..1..1    63    0       .1....11       6
33 = 1....1  :     .1..1.    21    0       .1.1.111       6 N
34 = 1...1.  :     1..1..    63    1       .11..111       6
35 = 1...11  :     ..1.11     9    0       .1..1..1       6 N
36 = 1..1..  :     .1.11.     7    0       ....11.1       3
37 = 1..1.1  :     1.11..    63    1       .11.11.1       6
38 = 1..11.  :     .11.11    63    0       .1.11.11       6
39 = 1..111  :     11.11.    21    1       .111.1.1       6 N
40 = 1.1...  :     1.1111    63    1       .11..111       6
41 = 1.1..1  :     .111.1    63    0       .1.11.11       6
42 = 1.1.1.  :     111.1.     3    1       .....111       2
43 = 1.1.11  :     11.111    63    1       .111..11       6
44 = 1.11..  :     1.11.1    63    1       .11.11.1       6
45 = 1.11.1  :     .11..1     7    0       ....1.11       3
46 = 1.111.  :     11..1.    63    1       .111..11       6
47 = 1.1111  :     1..111    63    1       .11....1       6
48 = 11....  :     ..11.1    21    0       .1.1.111       6 N
49 = 11...1  :     .11.1.     9    0       .1..1..1       6 N
50 = 11..1.  :     11.1..    63    1       .11.11.1       6
51 = 11..11  :     1.1.11    21    1       .111.1.1       6 N
52 = 11.1..  :     .1.1.1    63    0       .1.11.11       6
53 = 11.1.1  :     1.1.1.    63    1       .111..11       6
54 = 11.11.  :     .1.111     7    0       ....1.11       3
55 = 11.111  :     1.111.    63    1       .11....1       6
56 = 111...  :     .11111     9    0       .1..1..1       6 N
57 = 111..1  :     11111.    21    1       .111.1.1       6 N
58 = 111.1.  :     111111    63    1       .111..11       6
59 = 111.11  :     1111.1    63    1       .11....1       6
60 = 1111..  :     111..1    21    1       .111.1.1       6 N
61 = 1111.1  :     11...1    63    1       .11....1       6
62 = 11111.  :     1....1    63    1       .11....1       6
```

Figure 40.2-A: Minimal polynomials of the powers of a generator in GF(2^6). Polynomials of degree $n = 6$ that are non-primitive are marked with an 'N'. The trace of an element equals the coefficient of x^{n-1} of its minimal polynomial.

nomial can be determined from an irreducible polynomial C and a generator g modulo C by computing the minimal polynomial of g.

With a primitive polynomial (and the generator $g =$'x') the minimal polynomial of an element $f = g^k$ is primitive if k is a Lyndon word and $\gcd(k, n) = 1$. With a fast algorithm for the generation of Lyndon words one can therefore generate all primitive polynomials as shown in section 38.10 on page 849.

An alternative technique for the computation of the minimal polynomial of a given element is given in [238].

40.3 Fast computation of the trace vector

Let $C(x)$ be a polynomial with n roots a_0, a_1, ..., a_{n-1}

$$C(x) \;=\; \prod_{k=0}^{n-1} (x - a_k) \;=\; \sum_{k=0}^{n} c_k\, x^k \tag{40.3-1}$$

We define (following [257, sec.32])

$$s_k \;:=\; a_0^k + a_1^k + \ldots + a_{n-1}^k \tag{40.3-2}$$

Then, for $m = 1, \ldots, n$, we have *Newton's formula*:

$$m\, c_{n-m} \;=\; -\sum_{j=0}^{m-1} s_{m-j}\, c_{n-j} \tag{40.3-3}$$

Now let $C = c_0 + c_1\, x + c_2\, x^2 + \ldots + c_n\, x^n$ be an irreducible polynomial with coefficients in GF(p). Its roots are a, (and the conjugates) a^p, a^{p^2}, a^{p^3}, ..., $a^{p^{n-1}}$. Let $t_0 = n$ and $t_i = \mathrm{Tr}(a^i)$ (computationally x is a root of C, so $t_i = \mathrm{Tr}(x^i)$). Note that $t_0, \ldots, t_{n-1}$ are the elements of the trace vector, see relation 40.1-12 on page 877. Using $c_n = 1$ (monic polynomial C) and $t_j = s_j$ we rewrite Newton's formula as

$$\begin{aligned}
t_1 &= -1\, c_{n-1} && \text{(40.3-4a)}\\
t_2 &= -c_{n-1}\, t_1 - 2\, c_{n-2} && \text{(40.3-4b)}\\
t_3 &= -c_{n-1}\, t_2 - c_{n-2}\, t_1 - 3\, c_{n-3} && \text{(40.3-4c)}\\
t_4 &= -c_{n-1}\, t_3 - c_{n-2}\, t_2 - c_{n-3}\, t_1 - 4\, c_{n-4} && \text{(40.3-4d)}\\
t_5 &= -c_{n-1}\, t_4 - c_{n-2}\, t_3 - c_{n-3}\, t_2 - c_{n-4}\, t_1 - 5\, c_{n-5} && \text{(40.3-4e)}\\
&\;\;\vdots\\
t_k &= -c_{n-1}\, t_{k-1} - c_{n-2}\, t_{k-2} - \ldots - c_{n-k-1}\, t_1 - k\, c_{n-k} && \text{(40.3-4f)}
\end{aligned}$$

To compute the trace vector for the field GF(p^n), make the assignments in the given order, and finally compute $t_0 = n \bmod p$. The computation does not involve any polynomial modular reduction so the method can be worthwhile even for the determination of the trace of just one element.

With binary finite fields, the components with even subscripts can be computed as $t_{2k} = t_k$. During the computation we set $t_0 = 0$ and correct the value at the end of the routine. An implementation of the implied algorithm is [FXT: bpol/gf2n-trace.cc]:

```
1  ulong
2  gf2n_trace_vector_x(ulong c, ulong n)
3  // Return vector of traces of powers of x, where
4  //  x is a root of the irreducible polynomial C.
5  // Must have:  n == degree(C)
6  {
7      c &= ~( 2UL<<(n-1) );  // remove coefficient c[n]
```

```
 8
 9       ulong t = 1;  // set t[0]=1, will be corrected at the end
10       for (ulong k=1; k<n; ++k)
11       {
12           if ( k & 1 )  // k odd: use recursion
13           {
14               ulong cv = c >> (n-k);  // polynomials coefficients [n-1]..[n-k]
15               cv &= t;                // vector (j=1, k,  c[n-j]*t[k-j])
16               cv = parity(cv);        //    sum (j=1, k,  c[n-j]*t[k-j])
17               t |=  (cv<<k);
18           }
19           else  // k even:  copy t[k/2] to t[k]
20           {
21               t |= ( (t>>(k/2)) & 1 ) << k;
22           }
23       }
24
25       // correct t[0]:
26       t ^= ((n+1)&1);  // change low bit if n is even
27
28       return t;
29   }
```

The routine involves n computations of the parity. The complexity of the equivalent routine for large n has complexity $O(n^2)$ (n computations of sums with $\sim n$ summands).

For a binary polynomial C of odd degree n and all nonzero coefficients c_i at odd indices i we obtain $t_0 = 1$ and $t_i = 0$ for all $i \neq 0$, thereby the trace of any element is just the value of its lowest bit. In [47] it is shown that for $n \equiv \pm 3 \bmod 8$ the first nonzero coefficient c_k (with $k < n$) must appear at a position $k \geq n/3$.

With even degree and all nonzero odd coefficients c_i, c_j, c_j, ... at positions $i, j, k, \cdots < n/2$ the only nonzero components of the trace vector are t_{n-i}, t_{n-j}, t_{n-k}, ... Thereby polynomials of even degree with just one nonzero coefficient c_k where $k < n/2$ lead to only t_{n-k} being nonzero. A special case are trinomials $C = x^n + x^k + 1$ of even degree n and $k < n/2$ (k must be odd else C is reducible). The trace vector for all-ones polynomials ($C = \sum_{k=0}^{n} x^k$, see section 38.9.9 on page 846) the only zero component of the trace vector is t_0 (the degree must be even, else C is reducible). A detailed discussion of the properties of the trace vector is given in [5].

The following variant of the algorithm, suggested by Richard Brent [priv.comm.], shows that the computation is equivalent to a division of power series. Let R be the reciprocal polynomial of C, then (see [65, p.135])

$$\log\left(R(x)\right) \quad = \quad -\sum_{j=1}^{\infty} t_j \, x^j / j \tag{40.3-5}$$

Differentiating both sides gives

$$\frac{R'(x)}{R(x)} \quad = \quad -\sum_{j=1}^{\infty} t_j \, x^{j-1} \tag{40.3-6}$$

Using Newton's method for the inversion we obtain a computational cost of $\gamma\, M(n)$ where $M(n)$ is the cost for the multiplication of two power series up to order x^n and γ is a constant. The constant γ equals three if the division is performed by one inversion, which is two multiplications with the second order Newton iteration, and one final multiplication with $R'(x)$. For large n the multiplications should be done by either one of the splitting schemes suggested in [49] or by FFT methods such as given in [233]. Note added: relation 40.3-6 is given in [126, p.24].

40.4 Solving quadratic equations

We want to solve, in GF(2^n), the equation

$$a\,x^2 + b\,x + c \quad = \quad 0 \tag{40.4-1}$$

The fact that extracting a square root of an arbitrary element in GF(2^n) is easy does not enable us to solve the given equation. The formula $r_{0,1} = \left(-b \pm \sqrt{b^2 - 4ac}\right)/(2a)$ that works fine for real and complex numbers is of no help here: how should we divide by two?

Instead we transform the equation into a special form: divide by a: $x^2 + (b/a)\,x + (c/a) = 0$, substitute $x = z\,(b/a)$ to get $z^2\,(b/a)^2 + (b/a)^2\,z + (c/a) = 0$, and divide by $(b/a)^2$ to obtain

$$z^2 + z + C \quad = \quad 0 \quad \text{where} \quad C = \frac{a\,c}{b^2} \tag{40.4-2}$$

If r_0 is one solution of this equation then $r_1 = r_0 + 1$ is the other one: $z\,(z+1) = C$. The equation does not necessarily have a solution at all, the trace of C must be zero because we have $\mathrm{Tr}(C) = \mathrm{Tr}(z^2 + z) = \mathrm{Tr}(z^2) + \mathrm{Tr}(z) = \mathrm{Tr}(z) + \mathrm{Tr}(z) = 0$ for all $z \in \mathrm{GF}(2^n)$.

```
      n = 5  GF(2^n)
      c = 1..1.1 == x^5 + x^2 + 1  (polynomial modulus)
      mm= .11111   == 31 (prime)   (maximal order)
      h = .1.... (aux. bitmask)
      g = ....1. (element of maximal order)
      tv= ..1..1 (traces of x^i)
    tr1e= .....1 (element with trace=1)

       k:   f:=g**k  Tr(f)   RootOf(z^2+z=f)
       0:    ....1    1
       1:    ...1.    0      .1..1
       2:    ..1..    0      .1.11
       3:    .1...    1
       4:    1....    0      .1111
       5:    ..1.1    1
       6:    .1.1.    1
       7:    1.1..    0      ..1..
       8:    .11.1    0      11111
       9:    11.1.    1
      10:    1...1    1
      11:    ..111    1
      12:    .111.    1
      13:    1111.    1
      14:    111.1    0      1....
      15:    11111    0      11..1
      16:    11.11    0      1..1.
      17:    1..11    1
      18:    ...11    1
      19:    ..11.    0      ...1.
      20:    .11..    1
      21:    11...    1
      22:    1.1.1    1
      23:    .1111    0      1.11.
      24:    1111.    1
      25:    11..1    0      11.11
      26:    1.111    1
      27:    .1.11    0      111.1
      28:    1.11.    0      .11.1
      29:    .1..1    0      1.1..
      30:    1..1.    0      ..11.
```

Figure 40.4-A: Solutions of the equation $z^2 + z = f$ for all elements $f \in \mathrm{GF}(2^5)$ with trace zero.

The function [FXT: `gf2n_solve_quadratic()` in bpol/gf2n-solvequadratic.cc] transforms the equation into the reduced form:

```
1  bool gf2n_solve_quadratic(GF2n a, GF2n b, GF2n c, GF2n& r0, GF2n& r1)
2  // Solve  a*x^2 + b*x + c == 0
3  // Return  whether solutions exist.
4  // If so, the solutions are written to r0 and r1.
5  {
```

```
6       GF2n cc = a*c;
7       cc /= (b.sqr());  // cc = (a*c)/(b*b)
8       GF2n r;
9       bool q = gf2n_solve_reduced_quadratic(cc, r);
10      if ( !q )  return  false;
11
12      GF2n s = b / a;
13      r0 = r * s;
14      r1 = (r+GF2n::one) * s;
15
16      return  true;
17  }
```

The following function checks whether the reduced equation has a solution and if so, returns true and writes one solution to the variable r:

```
1   bool
2   gf2n_solve_reduced_quadratic(GF2n c, GF2n& r)
3   // Solve  z^2 + z == c
4   // Return  whether solutions exist.
5   // If so, one solutions is written to r.
6   // The other solution is r+1.
7   {
8       if ( 1==c.trace() )  return  false;
9
10      GF2n t( GF2n::tr1e );
11      GF2n z( GF2n::zero );
12      GF2n u( t );
13      for (ulong j=1; j<GF2n::n_; ++j)
14      {
15          GF2n u2 = u.sqr();
16          z = z.sqr();  z += u2*c;   // z = z*z + c*u*u
17          u = u2 + t;               // u = u*u + t
18      }
19      r = z;
20
21      return  true;
22  }
```

Figure 40.4-A shows the solutions to the reduced equations $x^2 + x = f$ for all elements f with trace zero [FXT: gf2n/gf2n-solvequadratic-demo.cc].

```
------- k=1: -------
    u=t^2 + t
    z=c*t^2
  z^2=c^2*t^4
z^2+z=c^2*t^4 + c*t^2

------- k=2: -------
    u=t^4 + t^2 + t
    z=(c^2 + c)*t^4 + c*t^2
  z^2=(c^4 + c^2)*t^8 + c^2*t^4
z^2+z=(c^4 + c^2)*t^8 + c*t^4 + c*t^2

------- k=3: -------
    u=t^8 + t^4 + t^2 + t
    z=(c^4 + c^2 + c)*t^8 + (c^2 + c)*t^4 + c*t^2
  z^2=(c^8 + c^4 + c^2)*t^16 + (c^4 + c^2)*t^8 + c^2*t^4
z^2+z=(c^8 + c^4 + c^2)*t^16 + c*t^8 + c*t^4 + c*t^2
     =(c^8 + c^4 + c^2)*t    + c*(t^8 + t^4 + t^2)    [using t^16=t]
     =( c + trace(c)  )*t    + c*( t + trace(t) )     [using x^8+x^4+x^2=x+trace(x)]
     =( c + 0  )*t           + c*( t + 1 )            [using trace(c)=0, trace(t)=1]
     =                         c                      [ z^2+z==c ]
```

Figure 40.4-B: Solving the reduced quadratic equation $z^2 + z = c$ in GF(2^4).

The implementation of the algorithm takes advantage of a precomputed element with trace one. At the

end of step $k \geq 1$ we have

$$u_k = \sum_{j=0}^{k} t^{2^k} \tag{40.4-3a}$$

$$z_k = \sum_{j=0}^{k-1} \left[t^{2^k} \sum_{i=0}^{k-1} c^{2^i} \right] \tag{40.4-3b}$$

Figure 40.4-B shows (for GF(2^4)) that this expression is the solution sought.

Routines for the solution of quadratic equations that do not depend on the class `GF2n` are given in [FXT: bpol/bitpolmod-solvequadratic.cc].

For GF(2^n) with n odd the solution of the reduced quadratic equation $z^2 + z = A$ can be computed via the *half-trace* of A which is defined as

$$H(A) = A + A^4 + A^{16} + \ldots + A^{4^{(n-1)/2}} \tag{40.4-4}$$

We have $H(A)^2 + H(A) = \mathrm{Tr}(A) + A$, so $H(A)$ is a solution of the reduced quadratic if $\mathrm{Tr}(A) = 0$. The half-trace of an element A in the field with field polynomial C can be computed via [FXT: `gf2n_half_trace()` in bpol/gf2n-trace.cc]:

```
1   ulong
2   gf2n_half_trace(ulong a, ulong c, ulong h)
3   {
4       ulong t = a;
5       ulong d = h;
6       while ( d>>=2 )
7       {
8           t = bitpolmod_square(t, c, h);
9           t = bitpolmod_square(t, c, h);
10          t ^= a;
11      }
12      return  t;
13  }
```

40.5 Representation by matrices *

```
n = 4  GF(2^n)
c = 1..11 == x^4 + x + 1  (polynomial modulus)

  k:  f:=g**k
  0:   ...1
  1:   ..1.
  2:   .1..        1...1..11.1.111
  3:   1...        .1..11.1.1111..
  4:   ..11        ..1..11.1.1111.
  5:   .11.        ...1..11.1.1111
  6:   11..
  7:   1.11
  8:   .1.1          M_0 =    M_1 =    M_2 =    M_3 =    M_4 =     M_5 =
  9:   1.1.          1...     ...1     ..1.     .1..     1..1      ..11
 10:   .111          .1..     1..1     ..11     .11.     11.1      1.1.
 11:   111.          ..1.     .1..     1..1     ..11     .11.      11.1
 12:   1111          ...1     ..1.     .1..     1..1     ..11      .11.
 13:   11.1
 14:   1..1
```

Figure 40.5-A: Powers of the primitive element in GF(2^4) with field polynomial $x^4 + x + 1$ (left), the list of powers rotated counter clockwise by 90 degree (top right), and the matrices obtained by taking 4 consecutive columns of the list (bottom right).

With a primitive polynomial modulus $C(x)$ a representation of the elements of GF(2^n) as matrices can be obtained from the powers of the primitive element ('x') in a surprisingly simply way: in the table of

powers of a generator take row k through $k+n-1$ to obtain the columns of matrices M_k, see figure 40.5-A. Now one has $M_k = M_1^k$ so one can use the matrices to represent the elements of GF(2^n).

The matrix $M := M_1$ is the *companion matrix* of the polynomial modulus g. The companion matrix of a polynomial $p(x) = x^n - \sum_{i=0}^{n-1} c_i\, x^i$ of degree n is defined as the $n \times n$ matrix

$$M \quad = \quad \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 & c_0 \\ 1 & 0 & 0 & \cdots & 0 & c_1 \\ 0 & 1 & 0 & \cdots & 0 & c_2 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & c_{n-2} \\ 0 & 0 & 0 & \cdots & 1 & c_{n-1} \end{bmatrix} \tag{40.5-1}$$

For polynomials $p(x) = \sum_{i=0}^{n} a_i\, x^i$ set $c_i := -a_i/a_n$.

The *characteristic polynomial* $c(x)$ of a $n \times n$ matrix M is defined as

$$c(x) \quad := \quad \det(x\, E_n - M) \tag{40.5-2}$$

where E_n is the $n \times n$ unit matrix. The roots of characteristic polynomial are the eigenvalues of the matrix. The characteristic polynomial of the companion matrix of a polynomial $p(x)$ equals $p(x)$. If $p(x)$ is the characteristic polynomial of a matrix M then $p(M) = 0$ (non-proof: set $x = M$ in relation 40.5-2, see [83] for a proof).

k:	[p_k(x)]^d	k:	[p_k(x)]^d
0:	[11]^4	7:	[1..11]^1
1:	[11..1]^1	8:	[11..1]^1
2:	[11..1]^1	9:	[11111]^1
3:	[11111]^1	10:	[111]^2
4:	[11..1]^1	11:	[1..11]^1
5:	[111]^2	12:	[11111]^1
6:	[11111]^1	13:	[1..11]^1
		14:	[1..11]^1

Figure 40.5-B: Characteristic polynomials of the powers of the generator x with the field GF(2^4) and the polynomial $x^4 + x + 1$.

Let $c_k(x)$ be the characteristic polynomial of the matrix $M_k = M^k$ and $p_k(x)$ the minimal polynomial of the element $g^k \in$ GF(2^n). Then

$$c_k(x) \quad = \quad [p_k(x)]^d \quad \text{where} \quad d = n/r \tag{40.5-3}$$

where r is the smallest positive integer so that $M_k^{2^r} = M_k$ For example, for the primitive modulus $C(x) = x^4 + x + 1$ (as above) the sequence of characteristic polynomials of the powers of the generator 'x' are shown in figure 40.5-B.

The trace of the matrix M^k is the d-th power of the *polynomial trace* of the minimal polynomial of g^k. The polynomial trace of $p(x) = x^n - (c_{n-1}\, x^{n-1} + \cdots + c_1\, x + c_0)$ equals c_{n-1} as can be seen from relation 40.5-1.

By construction, picking the first column of M_k gives the vector of the coefficients of the polynomial x^k modulo $C(x)$:

$$M^k\, [1, 0, 0, \ldots, 0]^T \quad \equiv \quad x^k \pmod{C(x)} \tag{40.5-4}$$

Finally, the characteristic polynomial of an element $a \in$ GF(2^n) in polynomial representation can be written as

$$p_a(x) \quad := \quad \prod_{k=0}^{n-1} \left(x - a^{2^k} \right) \tag{40.5-5}$$

Compare to relation 40.2-1 on page 882 for minimal polynomials.

40.6 Representation by normal bases

As a vector space over GF(2), we so far used the n basis vectors $x^0, x^1, x^2, x^3, \ldots, x^{n-1}$ to represent an element in GF(2^n). The arithmetic operations were the polynomial operations modulo an irreducible polynomial modulus C.

For certain irreducible polynomials it is possible to use the *normal basis* $x^1, x^2, x^4, x^8, \ldots, x^{2^{n-1}}$ to represent elements of GF(2^n). These polynomials are called *normal polynomials* or *N-polynomials*. To check whether a polynomial C is normal, compute $r_k = x^{2^k} \bmod C$ for $1 \le k \le n$, compute the nullspace of the matrix M whose k-th row is r_k. If the nullspace is empty (that is, $M \cdot v = 0$ implies $v = 0$) then the polynomial is normal.

The normality of a polynomial is equivalent to the fact that its roots are linearly independent (see section 38.5.1 on page 835 for the equivalence of computations modulo a polynomial and computations with linear combinations of its roots).

With a given polynomial basis the normal elements (sometimes called *free elements*) are the elements f of the field are those where $f^1, f^2, f^4, f^8, \ldots, f^{2^{n-1}}$ are linearly independent. The minimal polynomial of a normal element f is normal.

Addition and subtraction with a normal basis is again a simple XOR. Squaring an element can be achieved by a cyclic shift by one position. Note that $(x^{2^{n-1}})^2 = x^1$. Taking the square root is a cyclic shift in the other direction.

In normal basis representation the element one is the all-ones word. Thereby adding one is equivalent to complementing the binary word.

The trace can be computed easily with normal bases, it equals the parity of the binary word.

40.6.1 Multiplication and test for normality

Multiplication of two elements is achieved via a *multiplication matrix* M. Given two elements $a, b \in$ GF(2^n) in normal basis representation their product $p = a \cdot b$ is computed as follows: let a_k be the coefficient (zero or one) of x^{2^k} where $0 \le k \le n-1$. Then, for the first component p_0 of the product,

$$p_0 \quad = \quad a^T \cdot M \cdot b \tag{40.6-1}$$

and in general

$$p_k \quad = \quad \left(a^{-2^{k-1}}\right)^T \cdot M \cdot b^{-2^{k-1}} \tag{40.6-2}$$

That is, all components of the product are computed like the first but with a and b cyclically shifted.

A routine that checks whether a given polynomial c is normal and, if so, computes the multiplication matrix M, is [FXT: `bitpol_normal_q()` in bpol/bitpol-normal.cc] which proceeds as follows:

1. If the polynomial c is reducible, return false.
2. Compute the matrix A whose k-th row equals $x^{2^k} \bmod c$. If A is not invertible then (the nullspace is not empty and) c is not normal, so return false.
3. Set $D := A \cdot C^T \cdot A^{-1}$ where C is the companion matrix of c.
4. Compute the multiplication matrix M where $M_{i,j} := D_{j',i'}$, $i' := -i \bmod n$ and $j' := j - i \bmod n$. Return (true and) the matrix M.

Examples of the intermediate results for two different field polynomials are given in figure 40.6-A.

A C++ function implementing the multiplication algorithm is [FXT: `normal_mult()` in bpol/normal-mult.cc]:

```
Normal poly:  c=11111 =^= 4,3,2,1,0

    A=          A^-1=
      .1..        1111
      ..1.        1...
      1111        .1..
      ...1        ...1

    C^T=        D=A*C^T*A^{-1}=
      .1..        .1..
      ..1.        ...1
      ...1        1111
      1111        ..1.

    Multiplication matrix:  M=
      ..1.
      ..11
      11..
      .1.1
```

```
Normal poly:  c=111.11 =^= 5,4,2,1,0

A=          A^-1=
  .1...       11111
  ..1..       1....
  ....1       .1...
  .1.1.       1..1.
  1.111       ..1..

C^T=        D=A*C^T*A^{-1}=
  .1...       .1...
  ..1..       1..1.
  ...1.       ...11
  ....1       .11..
  111.1       ..1.1

Multiplication matrix:  M=
  .1...
  1..1.
  ...11
  .11..
  ..1.1
```

Figure 40.6-A: Matrices that occur with the computation of the multiplication matrix for the field polynomials $c = 1 + x + x^2 + x^3 + x^4$ (left) and $c = 1 + x + x^2 + x^4 + x^5$ (right).

```
1   ulong
2   normal_mult(ulong a, ulong b, const ulong *M, ulong n)
3   // Multiply two elements (a and b in GF(2^n))
4   //  in normal basis representation.
5   // The multiplication matrix has to be supplied in M.
6   {
7       ulong p = 0;
8       for (ulong k=0; k<n; ++k)
9       {
10          ulong v = bitmat_mult_Mv(M, n, b);
11          v = parity( v & a );  // dot product
12          p ^= ( v << k );
13          a = bit_rotate_right(a, 1, n);
14          b = bit_rotate_right(b, 1, n);
15      }
16      return  p;
17  }
```

We note that the algorithm is much more attractive for hardware implementations than for software, see [100].

40.6.2 Solving the reduced quadratic equation

The reduced quadratic equation $x^2 + x = f$ has two solutions if $\mathrm{Tr}(f) = 0$. If so, one solution $x = [x_0, x_1, \ldots, x_{n-1}]$ can be obtained as $x_k = \sum_{j=0}^{k} f_k$ where $f = [f_0, f_1, \ldots, f_{n-1}]$. This follows from observing that

$$x^2 + x \quad = \quad [x_0 + x_{n-1},\, x_0 + x_1,\, x_1 + x_2,\, \ldots,\, x_{n-2} + x_{n-1},\, x_{n-1} + x_0] \tag{40.6-3}$$

Now equate $x^2 + x = f$ and set $x_{n-1} = 0$ (setting $x_{n-1} = 1$ gives the complement which is also a solution). In C++ this translates to (see page 33) [FXT: `normal_solve_reduced_quadratic()` in bpol/normal-solvequadratic.cc]

```
1   ulong
2   normal_solve_reduced_quadratic(ulong c)
3   // Solve x^2 + x = c
4   // Must have:  trace(c)==0
5   // Return one solution x, the other solution equals 1+x,
6   // that is, the complement of x.
7   {
8       return  inverse_rev_gray_code(c);
9   }
```

```
Normal poly:  c=1111.1 == x^5 + x^4 + x^3 + x^2 + 1

   k =       :   f=g**k  Tr(f)   x^2+x==f
   0 = ..... :    11111    1
   1 = ....1 :    ....1    1
   2 = ...1. :    ...1.    1
   3 = ...11 :    .111.    1
   4 = ..1.. :    ..1..    1
   5 = ..1.1 :    1.111    .      x=.11.1
   6 = ..11. :    111..    1
   7 = ..111 :    .11.1    1
   8 = .1... :    .1...    1
   9 = .1..1 :    111.1    .      x=.1.11
  10 = .1.1. :    .1111    .      x=..1.1
  11 = .1.11 :    ..11.    .      x=...1.
  12 = .11.. :    11..1    1
  13 = .11.1 :    11...    .      x=.1...
  14 = .111. :    11.1.    1
  15 = .1111 :    1.1..    .      x=.11..
  16 = 1.... :    1....    1
  17 = 1...1 :    ..111    1
  18 = 1..1. :    11.11    .      x=.1..1
  19 = 1..11 :    1.11.    1
  20 = 1.1.. :    1111.    .      x=.1.1.
  21 = 1.1.1 :    ...11    .      x=....1
  22 = 1.11. :    .11..    .      x=..1..
  23 = 1.111 :    .1.1.    .      x=..11.
  24 = 11... :    1..11    1
  25 = 11..1 :    .1.11    1
  26 = 11.1. :    1...1    .      x=.1111
  27 = 11.11 :    ..1.1    .      x=...11
  28 = 111.. :    1.1.1    1
  29 = 111.1 :    1..1.    .      x=.111.
  30 = 1111. :    .1..1    .      x=..111
```

Figure 40.6-B: Solving the reduced quadratic equation $x^2 + x = f$ for powers $f = g^k$ of the generator $g = x$. The equation is solvable if the trace is zero, that is, the number of bits in the normal representation is even. The (primitive) field polynomial is $1 + x^2 + x^3 + x^4 + x^5$.

The highest bit of the result is zero exactly if the equation is solvable. The reversed Gray code is given in section 1.16.6 on page 45.

The program [FXT: gf2n/normalbasis-demo.cc] prints the powers of x in normal basis representation, see figure 40.6-B. By default a primitive normal polynomial from [FXT: bpol/normal-primpoly.cc] is used. If more than one argument is given the arguments should give the nonzero coefficients of the polynomial modulus.

40.6.3 The number of binary normal bases *

n :	A_n	n :	A_n	n :	A_n	n :	A_n
1:	1	11:	93	21:	27783	31:	28629151
2:	1	12:	128	22:	95232	32:	67108864
3:	1	13:	315	23:	182183	33:	97327197
4:	2	14:	448	24:	262144	34:	250675200
5:	3	15:	675	25:	629145	35:	352149525
6:	4	16:	2048	26:	1290240	36:	704643072
7:	7	17:	3825	27:	1835001	37:	1857283155
8:	16	18:	5376	28:	3670016	38:	3616800768
9:	21	19:	13797	29:	9256395	39:	5282242875
10:	48	20:	24576	30:	11059200	40:	12884901888

Figure 40.6-C: The number A_n of degree-n binary normal polynomials up to $n = 40$.

The number A_n of degree-n binary normal polynomials up to $n = 40$ is given in figure 40.6-C. A table of the values A_n for $1 \le n \le 130$ and their factorizations is given in [FXT: data/num-normalpoly.txt]. The

n:	B_n	n:	B_n	n:	B_n	n:	B_n
1:	1	11:	87	21:	23579	31:	28629151
2:	1	12:	52	22:	59986	32:	33552327
3:	1	13:	315	23:	178259	33:	78899078
4:	1	14:	291	24:	103680		
5:	3	15:	562	25:	607522		
6:	3	16:	1017	26:	859849		
7:	7	17:	3825	27:	1551227		
8:	7	18:	2870	28:	1815045		
9:	19	19:	13797	29:	9203747		
10:	29	20:	11255	30:	5505966		

Figure 40.6-D: The number B_n of degree-n binary primitive normal polynomials up to $n = 33$.

sequence A_n is sequence A027362 of [240].

The number B_n of degree-n binary primitive normal polynomials up to $n = 30$ is given in figure 40.6-D. This is sequence A107222 of [240].

40.6.3.1 Computation via exhaustive search

```
 1: c =  11..11...1  P ==   x^9 + x^8 + x^5 + x^4 + 1
 2: c =  11...1..11  P ==   x^9 + x^8 + x^4 + x + 1
 3: c =  11111...11  P ==   x^9 + x^8 + x^7 + x^6 + x^5 + x + 1
 4: c =  11.11.1.11  P ==   x^9 + x^8 + x^6 + x^5 + x^3 + x + 1
 5: c =  111....1.1  P ==   x^9 + x^8 + x^7 + x^2 + 1
 6: c =  111...1111  P ==   x^9 + x^8 + x^7 + x^3 + x^2 + x + 1
 7: c =  11.....1    ==     x^9 + x^8 + 1
 8: c =  11.111..11  P ==   x^9 + x^8 + x^6 + x^5 + x^4 + x + 1
 9: c =  1111.1.1.1  P ==   x^9 + x^8 + x^7 + x^6 + x^4 + x^2 + 1
10: c =  1111..1.11  P ==   x^9 + x^8 + x^7 + x^6 + x^3 + x + 1
11: c =  11.1.11.11  P ==   x^9 + x^8 + x^6 + x^4 + x^3 + x + 1
12: c =  11.1..1..1    ==   x^9 + x^8 + x^6 + x^3 + 1
13: c =  1111111.11  P ==   x^9 + x^8 + x^7 + x^6 + x^5 + x^4 + x^3 + x + 1
14: c =  1111...111  P ==   x^9 + x^8 + x^7 + x^6 + x^2 + x + 1
15: c =  11...1.1.1  P ==   x^9 + x^8 + x^4 + x^2 + 1
16: c =  11.1..1111  P ==   x^9 + x^8 + x^6 + x^3 + x^2 + x + 1
17: c =  11...11111  P ==   x^9 + x^8 + x^4 + x^3 + x^2 + x + 1
18: c =  111.111..1  P ==   x^9 + x^8 + x^7 + x^5 + x^4 + x^3 + 1
19: c =  1111.11..1  P ==   x^9 + x^8 + x^7 + x^6 + x^4 + x^3 + 1
20: c =  11..111.11  P ==   x^9 + x^8 + x^5 + x^4 + x^3 + x + 1
21: c =  11.11.11.1  P ==   x^9 + x^8 + x^6 + x^5 + x^3 + x^2 + 1
```

Figure 40.6-E: All normal binary polynomials of degree 9. Primitive polynomials are marked with '`P`'.

For small degrees all normal polynomials can be generated by selecting from the irreducible polynomials those that are normal. Using the mechanism that generate all irreducible polynomials via Lyndon words that is described in section 38.10 on page 849, the computation is a matter of minutes for $n < 25$.

The program [FXT: gf2n/all-normalpoly-demo.cc] prints all normal polynomials of a given degree n, its output for $n = 9$ is shown in figure 40.6-E

We can compute the number of normal (=: A_n) and primitive normal (=: B_n) binary polynomials for a small degrees n using that program. The table of the values B_n in figure 40.6-D was produced with the mentioned program, the computation up to $n = 30$ takes about 90 minutes. As noted in [132], no formula for the number of primitive normal polynomials is presently known. The proof that primitive normal bases exist for all finite extension fields was given 1987 in [183].

40.6.3.2 Cycles in the De Bruijn graph

Quite surprisingly, it turns out that A_n equals the number of cycles in the De Bruijn graph (see section 19.2.2 on page 377 and section 39.4 on page 862). Thereby for n a power of two the number A_n equals the number of binary De Bruijn sequences of length $2\,n$. No isomorphism between both objects (paths and polynomials) is presently known.

40.6.3.3 Invertible circulant matrices

```
      arg 1: 6 == n  [  n x n - matrices]  default=6
      arg 2: 2 == wh  [What to do: 0==>just count  1==>print words
                                   2==>also print matrix]  default=2

   v0 = 1..... [I]      v0 = 1.11.. [I]
         M =                  M =
        1.....               1.11..
        .1....               .1.11.
        ..1...               ..1.11
        ...1..               1..1.1
        ....1.               11..1.
        .....1               .11..1

   v0 = 111... [S]      v0 = 11111. [I]      v0 = 11.1.. [I]
         M =                  M =                  M =
        111...               11111.               11.1..
        .111..               .11111               .11.1.
        ..111.               1.1111               ..11.1
        ...111               11.111               1..11.
        1...11               111.11               .1..11
        11...1               1111.1               1.1..1

      n=6   #invertible=4   #singular=1
```

Figure 40.6-F: The length-6 Lyndon words of odd weight and the corresponding circulant matrices. Singular matrices are marked with '`[S]`', invertible matrices with '`[I]`'.

The number A_n of binary normal bases also equals the number of invertible circulant $n \times n$ matrices over GF(2). This is demonstrated in [FXT: gf2n/bitmat-circulant-demo.cc] whose output for $n = 6$ is shown in figure 40.6-F. The search uses the Lyndon words as periodic words would trivially lead to singular matrices. Further, Lyndon words with an even number of bits can be skipped as the vector $[1, 1, 1, \ldots, 1]$ is in the nullspace of the corresponding matrices.

If the set $\{\alpha, \alpha^2, \alpha^4, \alpha^8, \ldots, \alpha^{2^{n-1}}\}$ is a normal basis of GF(2^n) we say that α generates the normal basis. Considering the rows of a circulant matrix as some element β in a normal basis representation then the following rows are $\beta^2, \beta^4, \beta^8, \ldots, \beta^{2^{n-1}}$ and the matrix is invertible if β generates a normal basis. If α generates a normal basis then an element $\beta = \sum_{i=0}^{n-1} a_i\, \alpha^{2^i}$ generates a normal basis exactly if the polynomial $\sum_{i=0}^{n-1} a_i\, x^i$ is relatively prime to $x^n - 1$. Thereby, with a fast algorithm to generate Lyndon words, determine all elements that generate normal bases if one such element is known as follows: select the Lyndon words with an odd number of ones and test whether $\gcd(L(x), x^n - 1) = 1$ where $L(x)$ is the binary polynomial corresponding to the Lyndon word. If n is a power of two, then $x^n - 1 = (x - 1)^n$ and all Lyndon words with an odd number of ones are coprime to $x^n - 1$.

```
   L(x) = 111.11.  W(x) = 1....11       L(x) = 1.1.1..  W(x) = 1..1111
          M =            M^-1 =              M =            M^-1 =
         111.11.        1....11             1.1.1..        1..1111
         .111.11        11....1             .1.1.1.        11..111
         1.111.1        111...              ..1.1.1        111...11
         11.111.        .111...             1..1.1.        1111..1
         .11.111        ..111..             .1..1.1        11111..
         1.11.11        ...111.             1.1..1.        .11111.
         11.11.1        ....111             .1.1..1        ..11111
```

Figure 40.6-G: The inverse of a $n \times n$ circulant matrix over GF(2) can be found by computing the inverse $W(x)$ of its first row as a polynomial $L(x)$ modulo $x^n - 1$.

If the Lyndon word under consideration is taken as a polynomial $L(x)$ over GF(2) then the corresponding matrix is invertible exactly if $gcd(L(x), x^n - 1) = 1$. The first row inverse of a circulant matrix over GF(2) can be found by computing $W(x) = L(x)^{-1} \bmod x^n - 1$ where $L(x)$ is the binary polynomial with coefficients one where the Lyndon word has a one. As the inverse of a circulant matrix is also circulant, the remaining rows are cyclic shifts of $W(x)$. Two examples with $n = 7$ are shown in figure 40.6-G.

The equality of the number of invertible circulants and normal bases can also be seen as follows: choose a normal basis and test for each element f whether the set $f^1, f^2, f^4, f^8, \ldots, f^{2^{n-1}}$ is linearly independent. As squaring is a cyclic shift the matirces to be tested are the circulants we considered.

40.6.3.4 Factorization of $x^n - 1$

The factorization of the polynomial $x^n - 1$ over GF(2) can be used for the computation of A_n. The file [FXT: data/polfactdeg.txt] supplies the necessary information:

```
# Structure of the factorization of x^n-1 over GF(2):
 1:   [1]   [1*1]
 2:   [2]   [1*1]
 3:   [1]   [1*1 + 1*2]
 4:   [4]   [1*1]
 5:   [1]   [1*1 + 1*4]
 6:   [2]   [1*1 + 1*2]
 7:   [1]   [1*1 + 2*3]
 8:   [8]   [1*1]
 9:   [1]   [1*1 + 1*2 + 1*6]
10:   [2]   [1*1 + 1*4]
11:   [1]   [1*1 + 1*10]
12:   [4]   [1*1 + 1*2]
13:   [1]   [1*1 + 1*12]
14:   [2]   [1*1 + 2*3]
15:   [1]   [1*1 + 1*2 + 3*4]
16:  [16]   [1*1]
17:   [1]   [1*1 + 2*8]
```

An entry: `n: [e] [m1*d1 + m2*d2 + ... ]` says that $(x^n - 1) = P(x)^e$ and $P(x)$ factors into `m1` different irreducible polynomials of degree `d1`, `m2` different irreducible polynomials of degree `d2` and so on. As an example, for $n = 6$ we have

$$x^6 - 1 \;=\; \left[x^3 - 1\right]^2 \;=\; \left[(x+1)\left(x^2 + x + 1\right)\right]^2 \tag{40.6-4}$$

x^6 is the square ($e = 2$) of a product of one irreducible polynomial of degree one and one of degree two. Therefore we have the entry: `6: [2] [1*1 + 1*2]`. Another example, $n = 15$,

$$x^{15} - 1 \;=\; \left[(x+1)\left(x^2 + x + 1\right)\left(x^4 + x + 1\right)\left(x^4 + x^3 + 1\right)\left(x^4 + x^3 + x^2 + x + 1\right)\right]^1 \tag{40.6-5}$$

corresponding to the entry `15: [1] [1*1 + 1*2 + 3*4]`.

Now one has for the number of normal polynomials:

$$A_n \;=\; \frac{2^n}{n} \prod_i \left(1 - \frac{1}{2^{d_i}}\right)^{m_i} \tag{40.6-6}$$

Note that the quantity e does not appear in the formula. For example, with $n = 6$ and $n = 15$ we obtain

$$A_6 \;=\; \frac{2^6}{6} \cdot \left(1 - \frac{1}{2^1}\right)^1 \cdot \left(1 - \frac{1}{2^2}\right)^1 \;=\; \frac{64}{6} \cdot \frac{1}{2} \cdot \frac{3}{4} \;=\; 4 \tag{40.6-7a}$$

$$A_{15} \;=\; \frac{2^{15}}{15} \cdot \left(1 - \frac{1}{2^1}\right)^1 \cdot \left(1 - \frac{1}{2^2}\right)^1 \cdot \left(1 - \frac{1}{2^4}\right)^3 \;=\; 675 \tag{40.6-7b}$$

40.6.3.5 Efficient computation

It is actually possible to compute the number of degree-n normal binary polynomials without explicitly factorizing $x^n - 1$. We have

$$x^n - 1 \quad = \quad \prod_{d \backslash n} Y_d(x) \tag{40.6-8}$$

where $Y_d(x)$ is the d-th cyclotomic polynomial (see section 38.11 on page 850). We further know that $Y_d(x)$ factors into $\varphi(d)/r$ polynomials of degree r where $r = \mathrm{ord}_d(2)$ is the order of 2 modulo d. Let $a_n := A_n / \left(\frac{2^n}{n}\right)$, then a_n can for odd n be computed as

$$a_n \quad = \quad \prod_{d \backslash n} \left(1 - \frac{1}{2^r}\right)^{\varphi(d)/r} \tag{40.6-9}$$

The following pari/gp code works for any prime p:

```
1   p=2 /* global */
2   num_normal_p(n)=
3   {
4       local( r, i, pp );
5       pp = 1;
6       fordiv (n, d,
7           r = znorder(Mod(p,d));
8           i = eulerphi(d)/r;
9           pp *= (1 - 1/p^r)^i;
10      );
11      return( pp );
12  }
```

The number A_n can be computed (for arbitrary n) as $A_n = a_q \left(\frac{2^n}{n}\right)$ where q odd, and $n = q\,2^t$:

```
1   num_normal(n)=
2   {
3       local( t, q, pp );
4       t = 1;  q = n;
5       while ( 0==(q%p), q/=p; t+=1; );
6       /* here: n==q*p^t */
7       pp = num_normal_p(q);
8       pp *= p^n/n;
9       return( pp );
10  }
```

The quantity t is not used in the computation. The implementation is quite efficient: the computation of A_n for all $n \le 10,000$ takes less than three seconds. The computation of A_n for $n = 1234567 = 127 \cdot 9721$ (A_n is a number with 371,636 decimal digits) takes about 200 milliseconds.

40.6.4 Dual normal bases

Let $A = \{a_0, a_1, a_2, \cdots, a_{n-1}\}$ be a basis of GF(2^n). A basis $B = \{b_0, b_1, b_2, \cdots, b_{n-1}\}$ is said to be the *dual basis* (or *complementary basis*, or *trace-orthonormal basis*) of A if

$$\mathrm{Tr}\,(a_k\, b_j) \quad = \quad \delta_{k,j} \qquad \text{for} \quad 0 \le k, j < n \tag{40.6-10}$$

A basis that is its own dual is called *self-dual*. We treat only normal basis here. If α is a root of a normal polynomial C then $A = \{\alpha, \alpha^2, \alpha^4, \cdots, \alpha^{2^{n-1}}\}$ is a normal basis.

A necessary condition for C to be normal is that $\gcd(T, x^n - 1) = 1$ where

$$T \quad = \quad t_0 + t_1\, x + t_2\, x^2 + \ldots + t_{n-1}\, x^{n-1} \tag{40.6-11}$$

with $t_k = \mathrm{Tr}(\alpha\, \alpha^{2^k})$. The polynomial T can be computed via [FXT: bpol/normalpoly-dual.cc]

```
            C              T              C*       D = T^-1 (mod x^n-1)

    1:  11..11...1  P  .11....111     1111.1.1.1  .1.1111.11
    2:  11...1..11  P  .1..11..11     111....1.1  ..111111.1
    3:  11111...11  P  .1.1111.11     11..111.11  .11....111
    4:  11.11.1.11  P  ..1.11.1.1     1111111.11  .111..1111
    5:  111....1.1  P  ..111111.1     11...1..11  .1..11..11
    6:  111...1111  P  ..111111.1     11...1.1.1  .1..11..11
    7:  11.......1     .11....111     1111.11..1  .1.1111.11
    8:  11.111..11  P  ........1  S  11.111..11  ........1
    9:  1111.1.1.1  P  .1.1111.11     11..11...1  .11....111
   10:  1111..1.11  P  .111..1111     11.1.11.11  ..1.11.1.1
   11:  11.1.11.11  P  ..1.11.1.1     1111..1.11  .111..1111
   12:  11.1..1..1     ........1  S  11.1..1..1  ........1
   13:  1111111.11  P  .111..1111     11.11.1.11  ..1.11.1.1
   14:  1111...111  P  .111..1111     11.11.11.1  ..1.11.1.1
   15:  11...1.1.1  P  .1..11..11     111...1111  ..111111.1
   16:  11.1..1111  P  ........1  S  11.1..1111  ........1
   17:  11...11111  P  .1..11..11     111.111..1  ..111111.1
   18:  111.111..1  P  ..111111.1     11...11111  .1..11..11
   19:  1111.11..1  P  .1.1111.11     11.......1  .11....111
   20:  11..111.11  P  .11....111     11111...11  .1.1111.11
   21:  11.11.11.1  P  ..1.11.1.1     1111...111  .111..1111
```

Figure 40.6-H: All normal polynomials C of degree 9 and their polynomials T (left), their duals C^* and $D = T^{-1}$ (right). Primitive polynomials are marked with '`P`', self-dual $C = C^*$ are marked with '`S`'.

```
1   ulong
2   gf2n_ntrace(ulong c, ulong deg)
3   // Return vector V of traces v[k]=trace(ek), where
4   //   ek = x*x^(2^k), k=0..deg-1, and
5   //   x is a root of the irreducible polynomial C.
6   // Must have:  deg == degree(C)
7   {
8       if ( c==3 )  return 1UL;  // x+1 is self-dual
9
10      const ulong tv = gf2n_trace_vector_x(c, deg); // traces of x^k
11      ulong rt = 2UL;  // root of C
12      const ulong h = 1UL << (deg-1);  // aux
13
14      ulong v = 0;
15      for (ulong k=0; k<deg; ++k)
16      {
17          ulong ek = bitpolmod_times_x(rt, c, h);  // == x*x^(2^k)
18          ulong tk = gf2n_fast_trace(ek, tv);      // == sum(ek[i]*tk[i])
19          v |= (tk<<k);
20          rt = bitpolmod_square(rt, c, h);
21      }
22
23      return  v;
24  }
```

Now if C is normal then T has an inverse $D \equiv T^{-1} \bmod x^n - 1$. Write

$$D \;=\; d_0 + d_1\,x + d_2\,x^2 + \ldots + d_{n-1}\,x^{n-1} \tag{40.6-12a}$$

then

$$\beta \;=\; d_0\,\alpha + d_1\,\alpha^2 + d_2\,\alpha^4 + \ldots + d_{n-1}\,\alpha^{2^{n-1}} \tag{40.6-12b}$$

is a root of a normal polynomial so $B = \{\beta, \beta^2, \beta^4, \cdots, \beta^{2^{n-1}}\}$ is the dual (normal) basis of A. The following routine computes T, D, B, and finally the minimal polynomial C^* of β:

```
1   ulong
2   gf2n_dual_normal(ulong c, ulong deg, ulong ntc/*=0*/, ulong *ntd/*=0*/)
3   // Return the minimal polynomial CS for the dual (normal) basis
4   //  with the irreducible normal polynomial C.
5   // Return zero if C is not normal.
6   // Must have:  deg == degree(C).
7   // If ntc is supplied it must be equal to gf2n_ntrace(c, deg).
8   // If ntd is nonzero, ntc^-1 (mod x^deg-1) is written to it.
9   {
10      if ( 0==ntc )  ntc = gf2n_ntrace(c, deg);
```

```
11      const ulong d = bitpolmod_inverse(ntc, 1 | (1UL<<deg) );  // ntc=d^-1 (mod x^deg-1)
12      if ( 0==d )  return 0;  // C not normal
13      if ( 0!=ntd )  *ntd = d;
14
15      const ulong h = 1UL << (deg-1);  // aux
16      ulong alpha = 2UL;  // 'x', a root of C
17      ulong beta = 0;     // root of the dual polynomial
18      for (ulong m=d; m!=0; m>>=1)
19      {
20          if ( m & 1 )  beta ^= alpha;
21          alpha = bitpolmod_square(alpha, c, h);
22      }
23
24      ulong cs;  // minimal polynomial of beta
25      bitpolmod_minpoly(beta, c, deg, cs);
26
27      return cs;
28  }
```

Figure 40.6-H shows the normal bases of degree 9 and their duals, it was created with the program [FXT: gf2n/normalpoly-dual-demo.cc].

40.6.5 The number of self-dual normal basis

n:	S_n	n:	S_n	n:	S_n	n:	S_n	n:	S_n
1:	1	9:	3	17:	17	25:	205	33:	3267
2:	1	10:	4	18:	48	26:	320	34:	4352
3:	1	11:	3	19:	27	27:	513	35:	4095
4:	0	12:	0	20:	0	28:	0	36:	0
5:	1	13:	5	21:	63	29:	565	37:	7085
6:	2	14:	8	22:	96	30:	1920	38:	13824
7:	1	15:	15	23:	89	31:	961	39:	20475
8:	0	16:	0	24:	0	32:	0	40:	0

n:	Z_n	n:	Z_n	n:	Z_n	n:	Z_n	n:	Z_n
1:	0	9:	2	17:	17	25:	200	33:	2660
2:	1	10:	3	18:	25	26:	215	34:	2917
3:	1	11:	3	19:	27	27:	428	35:	
4:	0	12:	0	20:	0	28:	0	36:	
5:	1	13:	5	21:	57	29:	562	37:	
6:	1	14:	4	22:	60	30:	997	38:	
7:	1	15:	11	23:	87	31:	961	39:	
8:	0	16:	0	24:	0	32:	0	40:	

Figure 40.6-I: Number of self-dual normal basis (S_n, top) and self-dual normal basis where the field polynomial is primitive (Z_n, bottom).

Figure 40.6-I gives the number S_n of self-dual normal basis (top), and the number Z_n of such basis where the field polynomial is primitive (bottom). The field polynomial is the minimal polynomial of any of the basis elements. The sequence of values S_n is entry A135488 of [240], the values Z_n are entry A135498. No formula for the numbers Z_n is known, the values were computed with the program [FXT: gf2n/normalpoly-dual-demo.cc]. An expression for S_n is given in [155, theorem 5]. The following routine for computing the values S_n (for $p = 2$) is given by Max Alekseyev [priv.comm.]:

```
1  sdn(m,p) =
2  \\ Number of distinct self-dual normal bases of GF(p^m) over GF(p) where p is prime
3  {
4      local(F, f, g, s, c, d);
5      if ( p==2 && m%4==0, return(0) );
6
```

```
7       if ( !(m%p), /* p divides m */
8           s = m\p;
9           return( p^((p-1)*(s+(s*(p+1))%2)/2-1) * sdn(s,p) );
10      , /* else */
11          F = factormod( (x^m - 1)/(x - 1), p );
12          c = d = [];
13          for (i=1, matsize(F)[1],
14              f = lift(F[i,1]);
15              g = polrecip(f);
16              if ( f==g,  c = concat( c, vector(F[i,2],j,poldegree(f)/2) ); );
17              if ( lex(Vec(f), Vec(g))==1 ,
18                  d = concat( d, vector(F[i,2],j,poldegree(f)) );
19              );
20          );
21          return( 2^(p%2) * prod(i=1,#c, p^c[i] + 1) * prod(j=1,#d, p^d[j] - 1) / m );
22      );
23  }
```

We note that duality is defined for any basis, but no self-dual polynomial basis exists. See [128] and [184] for more information.

40.7 Conversion between normal and polynomial representation

If the field polynomial C is normal then conversion between the representations in polynomial and normal basis can be achieved as follows: Let Z be the $n \times n$ matrix whose k-th column equals $x^{2^k} \bmod C$ where n is the degree of C. If a is the polynomial representation then the normal representation is $b = Z^{-1} \cdot a$ (both a and b shall be column vectors).

```
k = bin(k):   f= g**k   == (normal)    trace(f)

0 = ....0 :      ...1  ==  11111       1
1 = ...10 :      ..1.  ==  ...1.       1
2 = ..1.. :      .1..  ==  ..1..       1
3 = ..110 :      .1...  ==  .111.       1
4 = .1... :      1....  ==  .1...       1
5 = .1.10 :      111.1  ==  1.111       .
6 = .110. :      .1110  ==  111..       1
7 = .1110 :      .111.  ==  .11.1       1            P2N=
8 = 1.... :      111..  ==  1....       1              11...
9 = 1..10 :      .1.10  ==  111.1       .              1.11.
10 = 1.1.. :      .1.1.  ==  .1111       .              1..11
11 = 1.110 :      1.1..  ==  ..11.       .              1..1.
12 = 1.10. :      1.1.1  ==  11..1       1              1....
13 = 1.110 :      1.110  ==  11...       .
14 = 1.110 :      1..11  ==  11.1.       1
15 = 1.111 :      11.11  ==  1.1..       .             N2P=
16 = 10... :      .1.11  ==  1....       1               ....1
17 = 10..1 :      1.11.  ==  ..111       1               1..1
18 = 10.1. :      1...1  ==  11.11       .               .1.1.
19 = 10.11 :      11111  ==  1.11.       1               ...11
20 = 1.1.. :      ...11  ==  1111.       .               ..11.
21 = 1.1.1 :      ..11.  ==  ...11       .
22 = 1.11. :      .11..  ==  .11..       .
23 = 1.111 :      11...  ==  .1.1.       .
24 = 110.. :      .11.1  ==  1..11       1
25 = 11..1 :      11.1.  ==  .1.11       1
26 = 11.1. :      .1..1  ==  1...1       .
27 = 11.11 :      1..1.  ==  ..1.1       .
28 = 111.. :      11..1  ==  1.1.1       1
29 = 111.1 :      .1111  ==  1..1.       .
30 = 1111. :      1111.  ==  .1..1       .
```

Figure 40.7-A: Conversion between normal- and polynomial representation with the (primitive) polynomial $c = 1 + x^2 + x^3 + x^4 + x^5$. The conversion matrices are given as `P2N` and `N2P`.

The implementation [FXT: `class GF2n` in bpol/gf2n.h] allows the conversion to the normal representation if the field polynomial is normal. In the initializer [FXT: `GF2n::init()` in bpol/gf2n.cc] the matrix Z (`n2p_tab[]`) and Z^{-1} (`p2n_tab[]`) are computed with the lines

```
1       // conversion to and from normal representation:
```

```
2       for (ulong k=0,s=2;  k<n_;  ++k)
3       {
4           n2p_tab[k] = s;
5           s = bitpolmod_square(s, c_, h_);
6       }
7       bitmat_transpose(n2p_tab, n_, n2p_tab);
8       is_normal_ = bitmat_inverse(n2p_tab, n_, p2n_tab);
```

The last line records whether the field polynomial is normal which is the case exactly if Z is invertible.

The functions [FXT: bpol/gf2n.cc]

```
1   ulong // static
2   GF2n::p2n(ulong f)
3   {  return  bitmat_mult_Mv(p2n_tab, n_, f);  }
4
5   ulong // static
6   GF2n::n2p(ulong f)
7   {  return  bitmat_mult_Mv(n2p_tab, n_, f);  }
```

allow conversions between the normal and polynomial representations. The method

```
ulong get_normal()  const  { return p2n(v_); }
```

provides a convenient way to obtain the normal representation of a given element.

This is demonstrated in [FXT: gf2n/gf2n-normal-demo.cc] where both the polynomial and the normal representation are given, see figure 40.7-A.

If the last argument of the initialization routine of the C++ class `GF2n`, `init(n, c, normalq)`, is set then a (primitive) normal polynomial will be used as field polynomial. A list of primitive normal polynomials is given in [FXT: bpol/normal-primpoly.cc].

40.8 Optimal normal bases (ONB)

The number of nonzero terms in the multiplication matrix determine the complexity (operation count) for the multiplication with normal bases. It turns out that for certain values of n there are normal bases of GF(2^n) whose multiplication matrices have at most two nonzero entries in each row (and column). Such bases are called *optimal normal bases* (ONB).

Optimal normal bases are especially interesting for hardware implementations because of both the highly regular structure of the multiplication algorithm and the minimal complexity with ONBs.

40.8.1 Type-1 optimal normal bases

A so-called *type-1 optimal normal basis* exists for n when $p := n + 1$ is prime and 2 is a primitive root modulo p (and for $n = 0$ and $n = 1$). The sequence of such n starts

```
0, 1, 2, 4, 10, 12, 18, 28, 36, 52, 58, 60, 66, 82, 100, 106, 130, 138, 148,
162, 172, 178, 180, 196, 210, 226, 268, 292, 316, 346, 348, 372, 378, 388,
418, 420, 442, 460, 466, 490, 508, ...
```

This is entry A071642 of [240]. One has always $n \equiv 2$ or $n \equiv 4$ modulo 8. A list of the corresponding primes is given in figure 39.6-B on page 866. The field polynomial corresponding to a type-1 ONB is the all-ones polynomial

$$c \;=\; \frac{x^p - 1}{x - 1} \;=\; 1 + x + x^2 + x^3 + \ldots + x^n \tag{40.8-1}$$

The order of these polynomials is $n + 1$ (for $n > 1$) so they are non-primitive for all $n \geq 3$.

The multiplication matrices are sparse: there is exactly one entry in the first row and column and exactly two in the other rows and columns. That is, the multiplication matrices for GF(n) with optimal normal

```
  Normal poly:  c=11111111111

   A:              A^-1:           C^T:          D=A*C^T*A^{-1}:  Mult. matrix M:
.1........      1111111111      .1........       .1.......          .....1....
..1.......      1.........      ..1.......       ........1.         .....11...
....1.....      .1........      ...1......       ....1.....         ...1...1..
........1.      ........1.      ....1.....       ......1...         ..1.....1.
.....1....      ..1.......      .....1....       ........1.         ......1.1
iiiiiiiiii      ....1.....      ......1...       iiiiiiiii1         ii........
.........1      .........1      .......1..       .....1....         .1......1.
...1....1..     ...1....1..     ........1.       ...1......         ..1.1.....
...1......      ...1......      .........1       ..1.......         ...1..1...
......1...      ......1...      iiiiiiiii1       .......1..         ....1....1
```

Figure 40.8-A: Matrices that occur with the computation of the multiplication matrix for the field polynomial $c = 1 + x + \ldots + x^{10}$.

basis have exactly $2n-1$ nonzero entries. For example, with $n = 10$ we obtain the matrix shown (together with the intermediate results) in figure 40.8-A. The equivalent data for $n = 4$ is shown in figure 40.6-A on page 891.

40.8.2 Type-2 optimal normal bases

```
 0:  -   1
 1: [1]  11
 2: [1]  111
 3: [1]  11.1  <--= x^4 + x^3 + 1
 4:  -   111.1
 5: [1]  11.111
 6: [1]  111..11
 7:  -   11.1...1
 8:  -   111.1...1
 9: [1]  11.111..11
10:  -   111..11.111
11: [1]  11.1...111.1
12:  -   111.1....11.1
13:  -   11.111.....111
14: [1]  111..11......11
15:  -   11.1...1.......1
16:  -   111.1...1.......1
17:  -   11.111..11......11
18: [1]  111..11.111.....111
19:  -   11.1...111.1....11.1
20:  -   111.1....11.1...111.1
21:  -   11.111.....111..11.111
22:  -   111..11......11.111..11
23: [1]  11.1...1.......111.1...1
24:  -   111.1...1........11.1...1
25:  -   11.111..11.........111..11
26: [1]  111..11.111..........11.111
27:  -   11.1...111.1...........111.1
28:  -   111.1....11.1............11.1
29: [1]  11.111.....111.............111
30: [1]  111..11......11..............11
31:  -   11.1...1.......1...............1
32:  -   111.1...1.......1...............1
33: [1]  11.111..11......11..............11
34:  -   111..11.111.....111.............111
35: [1]  11.1...111.1....11.1............11.1
36:  -   111.1....11.1...111.1...........111.1
37:  -   11.111.....111..11.111..........11.111
38:  -   111..11......11.111..11.........111..11
39: [1]  11.1...1.......111.1...1........11.1...1
40:  -   111.1...1........11.1...1.......111.1...1
```

Figure 40.8-B: The polynomials p_k up to $k = 40$ as binary strings. The entry in the second column is '`[1]`' the polynomial is irreducible (a field polynomial for a type-2 optimal normal basis).

A *type-2 optimal normal basis* exists for n if $p := 2\,n + 1$ is prime and either

- $n \equiv 1$ or $n \equiv 2$ modulo 4 and the order of 2 modulo p equals $2\,n$.
- $n \equiv 3$ mod 4 and the order of 2 modulo p equals n.

A type-2 basis exists for the following $n \leq 200$:

```
1, 2, 3, 5, 6, 9, 11, 14, 18, 23, 26, 29, 30, 33, 35, 39, 41, 50,
51, 53, 65, 69, 74, 81, 83, 86, 89, 90, 95, 98, 99, 105, 113, 119,
131, 134, 135, 146, 155, 158, 173, 174, 179, 183, 186, 189, 191, 194
```

This is entry A054639 of [240]. The corresponding normal polynomials p_n of degree n can be computed via the recurrence

$$p_0 := 1, \qquad p_1 := x+1 \tag{40.8-2a}$$

$$p_k := x\,p_{k-1} + p_{k-2} \tag{40.8-2b}$$

Compare to the recursion that transforms a linear hybrid cellular automaton into a binary polynomial relation 39.7-1 on page 868: the type-2 ONBs correspond to the most trivial LHCA defined by the rule having a single one as the lowest bit of the rule word.

As with type-1 ONBs, the multiplication matrices are sparse. The polynomials and multiplication matrices for $n = 6$ and $n = 9$ are

```
6,5,4,1,0      9,8,6,5,4,1,0
.1....         .1.....1...
1...1.         1...1.....
...11.         ....1..1..
..1..1         .......1.1
.11...         .11.....1.
...1.1         ...1..11..
               ..1.1.1...
               ...1....1.
```

The intermediate values with the computation of the multiplication matrix for $n = 5$ are shown in figure 40.6-A on page 891.

We note a relation of the polynomials p_k to the *Fibonacci polynomials* defined by

$$f_0 := 0, \qquad f_1 := 1 \tag{40.8-3a}$$

$$f_k := x\,f_{k-1} + f_{k-2} \tag{40.8-3b}$$

We have

$$p_k^2 = f_{2k+1} \tag{40.8-4}$$

The following relations hold over GF(2) (but not over $\mathbb{Z}$):

$$p_n(x) = \sum_{j=0}^{n} \binom{2n-j}{j} x^{n-j} = \sum_{j=0}^{n} \binom{n+j}{n-j} x^j \tag{40.8-5}$$

The binomial coefficient $\binom{n}{k}$ modulo 2 equals one if the binary expansion of k is a subset of the expansion of n. Using the trick from section 1.1.9 on page 6 we obtain a fast method for the computation of the polynomials p_n:

```
1   t2poly(n) = sum(j=0,n, (bitand(2*n-j, j)==j)*x^(n-j));
```

The value of binomial coefficients modulo a prime p can be computed via the relation

$$\binom{n}{k} \equiv \prod_{j=0}^{\infty} \left[\binom{n_j}{k_j} \right] \mod p \tag{40.8-6}$$

where $n = \sum_j n_j\,p^j$ and $k = \sum_j k_j\,p^j$ are the radix-p expansions (see the entry "Lucas Correspondence Theorem" in [268] and [116]). Moreover, the highest power of p that divides $\binom{n}{k}$ equals the number of carries when subtracting k from n in base p, see [198]. Especially if $k_j > n_j$ for any j then $\binom{n}{k} \equiv 0 \bmod p$. The computation above is obtained by setting $p = 2$.

The following relation is valid also over $\mathbb{Z}$:

$$p_n(x) = \sum_{j=0}^{n} \binom{n - \lfloor (j+1)/2 \rfloor}{\lfloor j/2 \rfloor} x^{n-j} \tag{40.8-7}$$

```
p7= 1*x^7 + 1*x^6 + 6*x^5 + 5*x^4 + 10*x^3 + 6*x^2 + 4*x + 1

    n \ k 0    1    2    3    4    5    6    7    8
    0:    1
    1:    1    1
    2:    1    2    1
    3:    1    3    3   *1
    4:    1    4   *6   *4    1
    5:    1   *5  *10   10    5    1
    6:   *1   *6   15   20   15    6    1
    7:   *1    7   21   35   35   21    7    1
    8:    1    8   28   56   70   56   28    8    1
```

Figure 40.8-C: Locations (starred entries) of the coefficients of the polynomial p_7 in Pascal's triangle.

The locations of the coefficients of the polynomials p_n in Pascal's triangle (see figure 6.0-A on page 175) lie on a rising diagonal. For $p_7(x) = 1\,x^7 + 1\,x^6 + 6\,x^5 + 5\,x^4 + 10\,x^3 + 6\,x^2 + 4\,x + 1$ they are shown in figure 40.8-C.

The sequence of values n such that a optimal normal basis over GF(2^n) exists is entry A136250 of [240]. The values up to 100 are

```
1, 2, 3, 4, 5, 6, 9, 10, 11, 12, 14, 18, 23, 26, 28, 29, 30, 33, 35, 36, 39, 41, 50,
51, 52, 53, 58, 60, 65, 66, 69, 74, 81, 82, 83, 86, 89, 90, 95, 98, 99, 100
```

40.9 Gaussian normal bases

The *type-t Gaussian normal basis* (GNB) generalize the optimal normal basis. The type-1 and type-2 GNBs are the corresponding ONBs. The multiplication matrices for type-t GNBs for $t > 2$ have more nonzero entries than the ONBs.

```
 1:  2, 4, 10, 12, 18, 28, 36, 52, 58, 60, 66, 82, 100, 106, 130
 2:  1, 2, 3, 5, 6, 9, 11, 14, 18, 23, 26, 29, 30, 33, 35
 3:  4, 6, 12, 14, 20, 22, 46, 52, 54, 60, 70, 76, 92, 94, 116
 4:  1, 3, 7, 9, 13, 15, 25, 37, 43, 45, 49, 67, 73, 79, 87
 5:  2, 12, 20, 26, 36, 42, 84, 92, 98, 108, 114, 132, 140, 164, 188
 6:  1, 2, 3, 5, 6, 7, 10, 11, 13, 17, 23, 26, 27, 30, 33
 7:  4, 28, 30, 54, 60, 70, 78, 94, 100, 108, 118, 126, 166, 196, 214
 8:  5, 9, 11, 17, 29, 35, 39, 51, 65, 71, 77, 95, 101, 107, 117
 9:  2, 4, 18, 20, 34, 42, 44, 58, 60, 68, 82, 84, 92, 98, 124
10:  1, 6, 7, 10, 13, 18, 19, 21, 27, 31, 42, 43, 46, 49, 54
```

Figure 40.9-A: For $t \in \{1, 2, \ldots, 10\}$: lowest values n such that a type-t GNB exists for GF(2^n).

A type-t GNB exists for n when $p := t\,n + 1$ is prime and $\gcd(n, t\,n/r_2) = 1$ where r_2 is the order of 2 modulo p. For n divisible by 8 no GNB exist. Figure 40.9-A shows, for $t \leq 10$, the first values n such that a type-t GNB exists. The sequences for $1 \leq t \leq 7$ are the following entries of [240]: A071642 (type-1), A054639 (type-2), A136415 (type-3), A137310 (type-4), A137311 (type-5), A137313 (type-6), A137314 (type-7), and A101284 (type-8). We implement the test using pari/gp:

```
1   gauss_test(n, t)=
2   { /* test whether a type-t Gaussian normal basis exists for GF(2^n) */
3       local( p, r2, g, d );
4       p = t*n + 1;
5       if ( !isprime(p), return( 0 ) );
6       if ( p<=2, return( 0 ) );
7       r2 = znorder( Mod(2, p) );
8       d = (t*n)/r2;
9       g = gcd(d, n);
10      return ( if ( 1==g, 1, 0) );
11  }
```

40.9.1 Computation of the multiplication matrix

An algorithm that computes the multiplication matrix for a type-t GNB proceeds as follows (The algorithm uses a vector $F[1, 2, \ldots, p-1]$):

1. Set $p = t\,n + 1$ (this is a prime), and compute an element r of order t modulo p.
2. For $k = 0, 1, \ldots, t-1$ do the following: set $j = r^k$ and for $i = 0, 1, \ldots, n-1$ set $F[j\,2^i] = i$.
3. Set the multiplication matrix M to zero.
4. For $i = 1, 2, \ldots, p-2$ add one to $M_{F[p-i],F[i+1]}$.
5. If t is odd set $h = n/2$ and do the following: for $i = 0, 1, \ldots, h-1$ increment $M_{i,h+i}$ and $M_{h+i,i}$.

Implementation in pari/gp:

```
gauss_nb(n, t)=
{ /* return multiplier matrix for type-t Gaussian normal basis */
  /* returned matrix is over Z and has to be multiplied by Mod(1,2) */
    local(p, r, F, w, x, nh, m, ir, ic);

    p = t*n + 1;
    r = znprimroot(p);  r = r^(n);  /* r has order t */
    F = vector(p-1);
    w = Mod(1, p);
    for (k=0, t-1,
        j = lift(w);
        for (i=0, n-1,
            F[j] = i;
            j+=j;  if (j>=p, j-=p);  /* 2*j mod p */
        );
        w *= r;
    );

    m = matrix(n, n);
    for (i=1, p-2,
        ir = F[p-i];  ic = F[i+1];
        m[ ir+1, ic+1 ] += 1;
    );

    if ( 1==(t%2),
        nh = n/2; /* odd t ==> even n */
        for (i=0, nh-1,
            ir = i;  ic = nh + i;
            ir += 1;  ic += 1;
            m[ir, ic] += 1;
            m[ic, ir] += 1;
        );
    );

    return ( m );
}
```

```
      n=7, t=4                          n=12, t=3

   M=          M mod 2=                   M=                M mod 2=
 .1..2..       .1.....              ....1.1...1.       ....1.1...1.
1.1..11        1.1..11              ...1.111....       ...1.111....
.1.111.        .1.111.              ....1...12...      ....1...1....
..12.1.        ..1..1.              .11..1...1..       .11..1...1..
2.1...1        ..1...1              1.......1.11.      1.......1.11.
.111..1        .111..1              .1.1.....1.1       .1.1.....1.1
.1..111        .1..111              11....2.....       11.........
                                    .11.1.....1.       .11.1.....1.
                                    ..2......1.1       .........1.1
                                    ...111..1...       ...111..1...
                                    1...1..1...1       1...1..1...1
                                    .....1..1.11       .....1..1.11
```

Figure 40.9-B: Multiplication matrices over $\mathbb{Z}$ and GF(2) for Gaussian normal bases with $n = 7$, $t = 4$ (left), and $n = 12$, $t = 3$ (right). Dots denote zeros.

The implementation computes M with entries in $\mathbb{Z}$ and has to be reduced modulo 2 before usage. Figure 40.9-B gives two examples.

The file [FXT: data/gauss-normalbasis.txt] lists for each $2 \le n \le 1032$ the smallest ten values of t so that there is a type-t GNB. Note that different values of t do not necessary lead to different multiplication matrices, especially for small values of n. For example, the modulo 2 reduced multiplication matrices for $n = 6$ and the 10 smallest values of t are:

```
 t=2:   t=3:   t=6:   t=10:   t=11:   t=23:   t=27:   t=30:   t=35:   t=55:
.1....  ..11.1 .1....  .1....  ..11.1 .1.11.  .1.11.  .1....  .1.11.  .1.11.
1...1.  ..11..  1.11.1  1.11.1  ..11..  1.111.  1.111.  1...1.  1.111.  1.111.
...11.  11..11 .1...1  .1...1  11..11  .1..1.  .1..1.  ..11..  .1..1.  .1..1.
..1..1  11....  .1..1.  .1..1.  11....  11....  11....  ..1..1  11....  11....
.11...  ..1..1  ...1.1  ...1.1  ..1..1  111..1  111..1  .11...  111..1  111..1
...1.1  1.1.11  .11.11  .11.11  1.1.11  ....11  ....11  ...1.1  ....11  ....11
```

40.9.2 Determination of the field polynomial

40.9.2.1 Algorithm with complex numbers

```
n=4 t=1:  p=5
  a(1)=2  w(1)=(-0.809016994374947 + 0.587785252292473 I)
  a(2)=4  w(2)=(+0.309016994374947 - 0.951056516295154 I)
  a(3)=3  w(3)=(-0.809016994374947 - 0.587785252292473 I)
  a(4)=1  w(4)=(+0.309016994374947 + 0.951056516295154 I)
z(x)=x^4 + x^3 + x^2 + x + 1
p(x)=x^4 + x^3 + x^2 + x + 1

n=4 t=3:  p=13
  a(1)=2  w(1)=(-1.15138781886600  + 1.72542218842201 I)
  a(2)=4  w(2)=(+0.651387818865997 - 0.522415803456408 I)
  a(3)=8  w(3)=(-1.15138781886600  - 1.72542218842201 I)
  a(4)=3  w(4)=(+0.651387818865997 + 0.522415803456408 I)
z(x)=x^4 + x^3 + 2*x^2 - 4*x + 3
p(x)=x^4 + x^3 + 1

n=4 t=7:  p=29
  a(1)=2  w(1)=(-1.59629120178363 - 0.509187583844044 I)
  a(2)=4  w(2)=(+1.09629120178363 + 2.64399848798351 I)
  a(3)=8  w(3)=(-1.59629120178363 + 0.509187583844044 I)
  a(4)=16 w(4)=(+1.09629120178363 - 2.64399848798351 I)
z(x)=x^4 + x^3 + 4*x^2 + 20*x + 23
p(x)=x^4 + x^3 + 1
```

Figure 40.9-C: Numerical values with the computation of the field polynomial for $n = 4$ and types $t \in \{1, 3, 7\}$. Note that the final result is identical for the types $t = 3$ and $t = 7$.

The normal polynomial corresponding to a type-t Gaussian basis can be computed as follows:

1. Set $p = t\,n + 1$ and determine r so that the order of r modulo p equals t.
2. For $1 \le k \le n$ compute $w_k = \sum_{j=0}^{t-1} \exp(a_k\, 2\,\pi\, i/p)$ where $a_k = 2^k\, r^j \bmod p$.
3. Let $z(x) = \prod_{k=1}^{n} (x - w_k)$, this is a polynomial with real integer coefficients.
4. Return the polynomial with coefficients reduced modulo 2.

The computation of the polynomial $z(x)$ uses complex (inexact) arithmetics. That its coefficients should be close to real integers can be used as a check.

The following pari/gp routine computes the complex polynomial. In order to keep the arguments for the exponential function small the values a_k are computed modulo p, the periodicity of $f(a) := \exp(a\, 2\,\pi\, i/p)$.

```
1  gauss_zpoly(n, t)=
2  { /* return field polynomial for type-t Gaussian normal basis
3       as polynomial over the complex numbers */
4
5     local(p, r, wk, tk1, tk, a, zp);
6     p = n*t + 1;
```

```
n=11 t=2:  p=23
  a(1)=2   w(1)=(+1.7088388090929771051 - 1.175494350822287508 E-38 I)
  a(2)=4   w(2)=(+0.9201300754623042520 + 5.877471754111437540 E-39 I)
  a(3)=8   w(3)=(-1.1533606442297342825 - 5.877471754111437540 E-39 I)
  a(4)=16  w(4)=(-0.6697592243419723039 - 2.938735877055718770 E-39 I)
  a(5)=9   w(5)=(-1.5514225814088396141 + 1.763241526233431262 E-38 I)
  a(6)=18  w(6)=(+0.4069120261052675797 + 1.193861450053885750 E-39 I)
  a(7)=13  w(7)=(-1.8344226030109060357 + 1.028557556969501569 E-38 I)
  a(8)=3   w(8)=(+1.3651062864373081657 + 1.175494350822287508 E-38 I)
  a(9)=6   w(9)=(-0.1364848267293419518 - 1.205340887073634651 E-39 I)
  a(10)=12 w(10)=(-1.9813718920726615047 + 2.277520304718182046 E-38 I)
  a(11)=1  w(11)=(+1.9258345746955985900 - 2.938735877055718770 E-38 I)

z(x)=x^11 + x^10 - 10*x^9 - 9*x^8 + 36*x^7 + 28*x^6 - 56*x^5 \
      - 35*x^4 + 35*x^3 + 15*x^2 - 6*x - 1
p(x)=x^11 + x^10 + x^8 + x^4 + x^3 + x^2 + 1

n=11 t=6:  p=67
  [--snip--]
z(x)=x^11 + x^10 - 30*x^9 - 63*x^8 + 220*x^7 + 698*x^6 - 101*x^5 \
      - 1960*x^4 - 1758*x^3 - 35*x^2 + 243*x - 29
p(x)=x^11 + x^10 + x^8 + x^5 + x^2 + x + 1

n=11 t=8:  p=89
  [--snip--]
z(x)=x^11 + x^10 - 40*x^9 - 19*x^8 + 482*x^7 + 84*x^6 - 2185*x^5 \
      + 102*x^4 + 3152*x^3 - 781*x^2 + 57*x - 1
p(x)=x^11 + x^10 + x^8 + x^5 + x^2 + x + 1

n=11 t=18:  p=199
  [--snip--]
z(x)=x^11 + x^10 - 90*x^9 - 115*x^8 + 2349*x^7 + 943*x^6 - 26327*x^5 \
      + 21284*x^4 + 102168*x^3 - 217794*x^2 + 148930*x - 30647
p(x)=x^11 + x^10 + x^8 + x^7 + x^6 + x^5 + 1
```

Figure 40.9-D: Numerical values with the computation of the field polynomial for $n = 11$ and types $t \in \{2, 6, 8, 18\}$. The final results are identical for $t = 6$ and $t = 8$.

```
7       r = znprimroot(p)^n;  \\ r has order t (mod p)
8
9       zp = 1;
10      tk1 = Mod(2,p);  tk = Mod(1,p);
11      for (k=1, n,
12          tk *= tk1;  \\ == Mod(2,p)^k;
13          wk = 0;
14          a = tk;
15          for (j=0, t-1,
16              wk += exp(2.0*I*Pi*lift(a)/p);
17              a *= r;
18          );
19          zp *= (x-wk);
20      );
21
22      return ( zp );
23  }
```

The final step uses pari/gp's function `round()` which rounds all coefficients of its polynomial argument:

```
1   gauss_poly(n, t)=
2   { /* return field polynomial for type-t Gaussian normal basis */
3       local(pp, zp);
4       zp = gauss_zpoly(n, t);
5       pp = round(real(zp));  /* rounds all coefficients */
6       pp *= Mod(1,2);  /* coefficients modulo 2 */
7       return( pp );
8   }
```

The results for type-1 bases can be verified using relation 40.8-1 on page 900, results with type-2 bases with relations 40.8-2a... 40.8-2b on page 902. The intermediate values occurring with the computation for $n = 4$ and the types $t \in \{1, 3, 7\}$ are shown in figure 40.9-C. The values for $n = 11$ and the types $t \in \{2, 6, 8, 18\}$ are shown in figure 40.9-D.

The computation can be optimized by using a trigonometric recursion as described in section 20.4.2 on page 402. We further exploit symmetry, and use real values if the type t is even:

```
vexp(p, t)=
{
    local( ve, ph, c, s, al, be,  cp, sp, tt );
    tt = 2.0*Pi/p;  \\ angle increment
    c = 1.0;  s = 0.0;  ga = ph;  al = 2.0*(sin(0.5*tt))^2; be = sin(tt);
    ve = vector(p);  ve[1] = 1.0;
    if ( t&1,  /* odd t, need complex values */
        for (j=1, (p-1)>>1,
            tt = c;
            c -= (al*tt+be*s);
            s -= (al*s -be*tt);
            ve[j+1]   = c + I*s;
            ve[p-j+1] = c - I*s;
        );
    , /* even t: can use real values */
        for (j=1, (p-1)>>1,
            tt = c;
            c -= (al*tt+be*s);
            s -= (al*s -be*tt);
            ve[j+1]   = c + s;
            ve[p-j+1] = c - s;
        );
    );
    return( ve );
}
```

The computation of the field polynomial needs two changes for even t:

```
gauss_vpoly(n, t)=
{
  [--snip--]
    ve = vexp(p, t);  \\ precompute trigonometric values
  [--snip--]
            wk += ve[lift(a)+1];  \\ was:  wk += exp(2.0*I*Pi*lift(a)/p);
  [--snip--]
}
```

For odd t we only need $n/2$ of the loop iterations:

```
    for (k=1, n\2,  \\ note: n/2 times
        tk *= tk1;  \\ == Mod(2,p)^k;
        wk = 0;
        a = tk;
        for (j=0, t-1,
            wk += ve[lift(a)+1];
            a *= r;
        );
        \\ use (x-(a+I*b))*(x-(a-I*b)) == x^2 - 2*a*x + (a^2 + b^2):
        zp *= (x^2-2*real(wk)*x+norm(wk));
    );
```

Note that the polynomial is always real throughout the computation.

40.9.2.2 Algorithm working in GF(2)

The following algorithm is a variation of what is given in [260].

1. Set $p = t\,n + 1$ and determine r so that the order of r modulo p equals t.
2. Set $M(x) = \sum_{k=0}^{p-1} x^k$. All computations are done modulo M.
3. If t equals 1 return M.
4. Set $F_0 = 1$ (modulo M).
5. For $1 \leq k \leq n$:
   (a) Set $Z_k = \sum_{j=0}^{t-1} x^{a(k,j)}$ (modulo M) where $a(k,j) = 2^k\, r^j \bmod p$.
   (b) Set $F_k = (x + Z_k)\; F_{i-1}$ (modulo M).

```
n=4  t=3:  p=13  \\ integer computation

   r= Mod(3, 13)  ord(r)=3 ==  t
   M= x^12 + x^11 + x^10 + x^9 + x^8 + x^7 + x^6 + x^5 + x^4 + x^3 + x^2 + x + 1

   ------- k=1
   Z= x^9 + x^3 + x
   F= x^9 + x^3 + 2*x
   ------- k=2
   Z= x^6 + x^5 + x^2
   F= x^11 + x^10 + x^9 + x^8 + 2*x^7 + 2*x^6 + x^5 + x^4 + 2*x^3 + 3*x^2 + x
   ------- k=3
   Z= -x^11 - x^9 - x^8 - x^7 - x^6 - x^5 - x^3 - x^2 - x - 1
   F= -x^11 - x^10 - 2*x^9 + x^7 + 2*x^6 + x^3 - 1
   ------- k=4
   Z= x^11 + x^8 + x^7
   F= x^4 - x^3 + 2*x^2 + 4*x + 3

 ==> x^4 - x^3 + 2*x^2 + 4*x + 3  ==   x^4 + x^3 + 1 (mod 2)
```

```
n=4  t=3:  p=13  \\ computation over GF(2)

   r= Mod(3, 13)  ord(r)=3 ==  t
   M= x^12 + x^11 + x^10 + x^9 + x^8 + x^7 + x^6 + x^5 + x^4 + x^3 + x^2 + x + 1

   ------- k=1
   Z= x^9 + x^3 + x
   F= x^9 + x^3 + 2*x
   ------- k=2
   Z= x^6 + x^5 + x^2
   F= x^11 + x^10 + x^9 + x^8 + x^5 + x^4 + x^2 + x
   ------- k=3
   Z= x^11 + x^9 + x^8 + x^7 + x^6 + x^5 + x^3 + x^2 + x + 1
   F= x^11 + x^10 + x^7 + x^3 + 1
   ------- k=4
   Z= x^11 + x^8 + x^7
   F= x^4 + x^3 + 1

 ==>  x^4 + x^3 + 1
```

Figure 40.9-E: Computation of the field polynomial for $n = 4$ and $t = 3$ with polynomials over the integers (top) and polynomials over GF(2) (bottom).

6. Return F_n with coefficients reduced modulo 2.

The intermediate quantities in the computation for $n = 4$ and $t = 3$ are shown at the top of figure 40.9-E. The result is a polynomial over the integers identical to the one computed with the algorithm that uses complex numbers. When all polynomials are taken over GF(2) the computation proceeds as shown at the bottom of figure 40.9-E. Implementation in pari/gp:

```
1   gauss_poly2(n, t)=
2   {  /* return field polynomial for type-t Gaussian normal basis */
3       local(p, M, r, F, t21, t2, Z);
4
5       p = t*n + 1;
6       r = znprimroot(p)^n;  \\ element of order t mod p
7       M = sum(k=0, p-1, 'x^k);  \\ The polynomial modulus
8       M *= Mod(1,2);  \\ ... over GF(2)
9       if ( 1==t,  return( M ) );  \\ for type 1
10
11      F = Mod(1, M);
12      t21 = Mod(2,p);  t2 = Mod(1,p);
13      for (k=1, n,
14          Z = sum(j=0, t-1, Mod('x^lift(t2*r^j), M) );
15          F = ('x+Z)*F;
16          t2 *= t21;
17      );
18      return ( lift(F) );
19  }
```

While the algorithm avoids inexact arithmetic the polynomial modulus M is of degree $p - 1 = n\,t$ which can be large for large t. In practice the computation with complex numbers is much faster. It finishes in less than a second for $n = 620$ and $t = 3$ (and a working precision of 150 decimal digits) while the exact method needs about two minutes.

Using the redundant modulus $x^p - 1 = (x-1) \cdot M$ gives a significant speedup:

```
gauss_poly2(n, t)=
{
    local(p, M, r, F, t21, t2, Z);
    p = t*n + 1;
    r = znprimroot(p)^n;  \\ element of order t mod p
    M = 'x^p - 1;  \\ Use redundant modulus (instead of sum(k=0, p-1, 'x^k))
    M *= Mod(1,2);  \\ ... over GF(2)

    if ( 1==t,  return( sum(k=0, p-1, 'x^k) ) );  \\ for type 1

  [--snip--]   \\ main loop as before

    \\ final reduction for redundant modulus:
    M = sum(k=0, p-1, 'x^k);  \\ The polynomial modulus
    F = lift( Mod( lift(F), M) );
    return ( F );
}
```

Now computation of the polynomial for $n = 620$ and $t = 3$ takes less than 9 seconds. The final reduction can be simplified by observing that no reduction is needed if the constant coefficient is one, and else all coefficients just have to be negated. So the end of the routine can be changed to

```
    \\ final reduction for redundant modulus (simplified):
    F = lift(F);
    if ( 0==polcoeff(F,0), F=sum(k=0, n, (1-polcoeff(F,k))*'x^k) );
    return ( F );
}
```

40.9.2.3 Computation using the multiplication matrix

For even t the polynomials can also be determined by the relation

$$F(x) \quad = \quad \det\left(x \ \mathrm{id} - M\right) \tag{40.9-1}$$

where M is the multiplication matrix and id is the $n \times n$ unit matrix. This leads to the one-liner

```
gauss_mpoly(n,t)= return ( charpoly(gauss_nb(n,t) )
```

The method is competitive only when t is large.

Appendix A

The electronic version of the book

How to make the hyperlinks work

The hyperlink showing as [FXT: bits/revbin.h] points to `file:~/.fxtdir/bits/revbin.h`. To make this work on your machine you may want to create (in your home directory) a soft-link to the directory of the FXTsources. For example, assuming that the package is located at `~/work/fxt`, execute the following statement:

```
ln -sv ~/work/fxt ~/.fxtdir
```

Similarly, for hfloat, do

```
ln -sv ~/work/hfloat ~/.hfloatdir
```

Test with the hyperlink [hfloat: src/hf/funcsrt.cc] which points to
`file:~/.hfloatdir/src/hf/funcsrt.cc`.

You may want to add the following lines to your `file:~/.mailcap`:

```
text/plain;/usr/bin/emacs -no-site-file %s &
text/x-csrc;/usr/bin/emacs -no-site-file %s &
text/x-chdr;/usr/bin/emacs -no-site-file %s &
text/x-c++src;/usr/bin/emacs -no-site-file %s &
text/x-c++hdr;/usr/bin/emacs -no-site-file %s &
```

Here the editor emacs is used for viewing plain text, C and C++ sources and headers.

Appendix B

Machine used for benchmarking

The machine used for performance measurements is a AMD64 (Athlon64) clocked at 2.2 GHz with dual channel double data rate (DDR) clocked at 200 MHz ('800 MHz'). It has 512 kB (16-way associative) second level cache and separate first level caches for data and instructions, each 64 kB (and 2-way associative). Cache lines are 64 bytes (8 words, 512 bits). The memory controller is integrated in the CPU.

The CPU has 16 general purpose (64 bit) registers that are addressable as byte, 16 bit word, 32 bit word, or 64 bit (full) word. These are used for integer operations and for passing integer function arguments. There are 16 (128 bit, SSE) registers that are used for floating point operations and for passing floating point function arguments. The SSE registers are SIMD registers. Additionally, there are 8 (legacy, x87) FPU registers.

The performance-wise interesting information reported by the CPUID instruction is:

```
Vendor: AuthenticAMD
Name: AMD Athlon(tm) 64 Processor 3400+
  Family: 15,  Model: 15,  Stepping: 0
Level 1 cache (data):  64 kB,  2-way associative.
  64 bytes per line,  lines per tag: 1.
Level 1 cache (instr):  64 kB,  2-way associative.
  64 bytes per line,  lines per tag: 1.
Level 2 cache:  512 kB,  16-way associative
  64 bytes per line,  lines per tag: 1.

Max virtual addr width: 48
Max physical addr width: 40
Features:
  lm: Long Mode (64-bit mode)
  mtrr: Memory Type Range Registers
  tsc: Time Stamp Counter
  fpu: x87 FPU
  3dnow: AMD 3DNow! instructions
  3dnowext: AMD Extensions to 3DNow!
  mmx: Multimedia Extensions
  mmxext: AMD Extensions to MMX
  sse: Streaming SIMD Extensions
  sse2: Streaming SIMD Extensions-2
  cmov: CMOV instruction (plus FPU FCMOVCC and FCOMI)
  cx8: CMPXCHG8 instruction
  clflush: CLFLUSH instruction
  fxsr: FXSAVE and FXRSTOR instructions
```

Special instructions as SIMD, prefetch and non-temporal moves are not used unless explicitly noted.

See [135] for a comparison of instruction latencies and throughput for various x86 CPU cores. You do want to study the cited document *before* buying an x86-based system.

The compiler used was the GNU (C and C++) compiler [119].

Appendix C

The pseudo language Sprache

Many algorithms in this book are given in a pseudo language called Sprache. Sprache is meant to be immediately understandable for everyone who ever had contact with programming languages like C, FORTRAN, Pascal or Algol. Sprache is hopefully self explanatory. The intention of using Sprache instead of completely relying an mathematical formulas (like tensor formalism) or algorithm description by words is to minimize the work it takes to translate the given algorithm to one's favorite programming language. It should be mere syntax adaptation.

By the way, 'Sprache' is the German word for language.

```
// a comment:
// comments are useful.

// assignment:
t := 2.71

// parallel assignment:
{s, t, u} := {5, 6, 7}
// same as:
s := 5
t := 6
u := 7

{s, t} := {s+t, s-t}
// same as (avoiding the temporary):
temp := s + t
t := s - t
s := temp

// if conditional:
if a==b then  a:=3

// with block
if a>=3 then
{
    // do something ...
}

// a function returns a value:
function plus_three(x)
{
    return  x + 3
}

// a procedure works on data:
procedure increment_copy(f[],g[],n)
// real f[0..n-1] input
// real g[0..n-1] result
{
    for k:=0 to n-1
    {
        g[k] := f[k] + 1
    }
}

// for loop with stepsize:
```

```
for i:=0 to n step 2  // i:=0,2,4,6,...
{
    // do something
}

// for loop with multiplication:
for i:=1 to 32 mul_step 2
{
    print i, ", "
}
```

will print `1, 2, 4, 8, 16, 32,`

```
// for loop with division:
for i:=32 to 8 div_step 2
{
    print i, ", "
}
```

will print `32, 16, 8,`

```
// while loop:
i:=5
while i>0
{
    // do something 5 times...
    i := i - 1
}
```

The usage of `foreach` emphasizes that no particular order is needed in the array access (so parallelization is possible):

```
procedure has_element(f[],x)
{
    foreach t in f[]
    {
        if t==x then return TRUE
    }
    return FALSE
}
```

Emphasize type and range of arrays:

```
real     a[0..n-1],    // has n elements (floating point reals)
complex  b[0..2**n-1]  // has 2**n elements (floating point complex)
mod_type m[729..1728]  // has 1000 elements (modular integers)
integer  i[]           // has ? elements (integers)
```

Arithmetical operators: `+, -, *, /, %` and `**` for powering. Arithmetical functions: `min(), max(), gcd(), lcm(), ...`

Mathematical functions: `sqr(), sqrt(), pow(), exp(), log(), sin(), cos(), tan(), asin(), acos(), atan(), ...`

Bitwise operators: `~, &, |, ^` for bit-wise complement, AND, OR, XOR, respectively. Bit shift operators: `A<<3` shifts (the integer) `A` 3 bits to the left `A>>1` shifts `A` 1 bits to the right.

Comparison operators: `==, !=, <, > ,<=, >=`

There is no operator '`=`' in Sprache, only '`==`' (for testing equality) and '`:=`' (assignment operator).

A well-known constant: `PI` $= 3.14159265\ldots$

The complex square root of minus one in the upper half plane: `I` $= \sqrt{-1}$

Boolean values: `TRUE` and `FALSE`

Logical operators: `NOT, AND, OR, XOR`

Modular arithmetic: `x := a * b mod m` shall do what it says, `i := a**(-1) mod m` shall set `i` to the modular inverse of `a`.

Appendix D

The pari/gp language

We give a short introduction to the pari/gp language.

From the manual page of pari [210] (slightly edited):

```
NAME       gp - PARI calculator

SYNOPSIS   gp [-emacs] [-f] [-test] [-q] [-s stacksize] [-p primelimit]

DESCRIPTION
       Invokes the PARI-GP calculator. This is an advanced programmable calcu-
       lator, which computes symbolically as  long  as  possible,  numerically
       where  needed,  and  contains  a  wealth  of number-theoretic functions
       (elliptic curves, class field theory...).  Its basic data types are

       integers, real numbers, exact rational numbers,  algebraic  numbers,
       p-adic numbers, complex numbers,
       modular integers,
       polynomials and rational functions,
       power series,
       binary quadratic forms,
       matrices, vectors, lists,
       character strings,
       and recursive combinations of these.
```

Interactive usage

To use pari/gp interactively, just type `gp` at your command line prompt. A startup message like the following will appear:

```
                          GP/PARI CALCULATOR Version 2.3.3 (released)
                        amd64 running linux (x86-64 kernel) 64-bit version
                          compiled: Jan 15 2008, gcc-4.2.1 (SUSE Linux)
                         (readline v5.2 enabled, extended help available)

                              Copyright (C) 2000-2006 The PARI Group

PARI/GP is free software, covered by the GNU General Public License, and comes WITHOUT ANY
WARRANTY WHATSOEVER.

Type ? for help, \q to quit.
Type ?12 for how to get moral (and possibly technical) support.

parisize = 8000000, primelimit = 500000
?
```

The question mark in the last line is a prompt, the program is waiting for your input.

```
? 1+1
%1 = 2
```

Here we successfully computed one plus one. Next we compute a factorial:

```
? 44!
%2 = 265827157478844876804362581101461589031963852800000000
```

Integers are of unlimited precision, the practical limit is the amount of physical RAM. For floating point numbers, the precision (number of decimal digits) can be set as follows

```
? default(realprecision,55)
%3 = 55
? sin(1.5)
%4 = 0.99749498660405443094172337114148732270665142592211582 19
```

The history numbers `%N` (where `N` is a number) can be used to recall the result of a prior computation:

```
? %4
%5 = 0.99749498660405443094172337114148732270665142592211582 19
```

The output of the result of a calculation can suppressed using a semicolon at the end of the command. This can be useful for timing purposes:

```
? default(realprecision,10000)
%5 = 10000
? sin(2.5);
? ##
  ***   last result computed in 100 ms.
```

The command `##` gives the time used for the last computation.

The printing format can be set independently of the precision used:

```
? default(realprecision,10000);
? default(format,"g.15");
? sin(2.5)
%6 = 0.598472144103956
```

Command line completion is available, typing `si`, then the tab-key, gives a list of builtin functions whose names start with `si`:

```
? si
  sigma      sign       simplify   sin        sinh       sizebyte   sizedigit
```

You can get the help text by using the question mark, followed by the help topic:

```
? ?sinh
  sinh(x): hyperbolic sine of x.
```

A help overview is invoked by a single question mark

```
? ?
Help topics: for a list of relevant subtopics, type ?n for n in
  0: user-defined identifiers (variable, alias, function)
  1: Standard monadic or dyadic OPERATORS
  2: CONVERSIONS and similar elementary functions
  3: TRANSCENDENTAL functions
  4: NUMBER THEORETICAL functions
  5: Functions related to ELLIPTIC CURVES
  6: Functions related to general NUMBER FIELDS
  7: POLYNOMIALS and power series
  8: Vectors, matrices, LINEAR ALGEBRA and sets
  9: SUMS, products, integrals and similar functions
 10: GRAPHIC functions
 11: PROGRAMMING under GP
 12: The PARI community
```

Select a section by its number:

```
? ?7
  O                 deriv             eval              factorpadic
  intformal         padicappr         polcoeff          polcyclo
  poldegree         poldisc           poldiscreduced    polhensellift
  polinterpolate    polisirreducible  pollead           pollegendre
  polrecip          polresultant      polroots          polrootsmod
  polrootspadic     polsturm          polsubcyclo       polsylvestermatrix
  polsym            poltchebi         polzagier         serconvol
  serlaplace        serreverse        subst             substpol
  substvec          taylor            thue              thueinit
```

You should try both of the following

```
? ??tutorial
  displaying 'tutorial.dvi'.
? ??
  displaying 'users.dvi'.
```

A short overview (which you may want to print) of most functions can be obtained via

```
? ??refcard
  displaying 'refcard.dvi'.
```

A session can be ended by either entering `quit` or just hitting control-d.

Builtin operators and basic functions

There are the 'usual' operators `+`, `-`, `*`, `/`, `^` (powering), and `%` (modulo). The operator `\` gives the integer quotient without remainder. The assignment operator is `=`. C-style shortcuts are available, for example `t+=3` is the same as `t=t+3`.

The increment by one can be abbreviated as `t++`, the decrement as `t--`. [Technical note: these behave as the C-language pre-increment (and pre-decrement), that is, the expression evaluates to `t+1`, not `t`. There is no post-increment or post-decrement in pari/gp.]

Comparison operators are `==`, `!=` (alternatively `<>`), `>`, `>=`, `<`, and `<=`. Logical operators are `&&` (and), (or), and `!` (not)

Bit-wise operations for integers are

```
  bitand  bitneg  bitnegimply  bitor  bittest  bitxor
```

and

```
shift(x,n): shift x left n bits if n>=0, right -n bits if n<0.
shiftmul(x,n): multiply x by 2^n (n>=0 or n<0)
```

One can also use the operators `>>` and `<<`, as in the C-language, and the shortcuts `>>=` and `<<=`.

An overview of basic functions is obtained as

```
? ?2
  Col          List         Mat          Mod          Pol          Polrev       Qfb
  Ser          Set          Str          Strchr       Strexpand    Strtex       Vec
  Vecsmall     binary       bitand       bitneg       bitnegimply  bitor        bittest
  bitxor       ceil         centerlift   changevar    component    conj         conjvec
  denominator  floor        frac         imag         length       lift         norm
  norml2       numerator    numtoperm    padicprec    permtonum    precision    random
  real         round        simplify     sizebyte     sizedigit    truncate     valuation
  variable
```

Here are a few:

```
sign(x): sign of x, of type integer, real or fraction.
max(x,y): maximum of x and y.
min(x,y): minimum of x and y.
abs(x): absolute value (or modulus) of x.

floor(x): floor of x = largest integer<=x.
ceil(x): ceiling of x=smallest integer>=x.
frac(x): fractional part of x = x-floor(x)
```

An overview of sums, products, and some numerical functions:

```
? ?9
  intcirc            intfouriercos      intfourierexp      intfouriersin
  intfuncinit        intlaplaceinv      intmellininv       intmellininvshort
  intnum             intnuminit         intnuminitgen      intnumromb
  intnumstep         prod               prodeuler          prodinf
  solve              sum                sumalt             sumdiv
  suminf             sumnum             sumnumalt          sumnuminit
  sumpos
```

For example:

```
sum(X=a,b,expr,{x=0}): x plus the sum (X goes from a to b) of expression expr.
prod(X=a,b,expr,{x=1}): x times the product (X runs from a to b) of expression.
```

Basic data types

Strings:

```
? a="good day!"
  "good day!"
```

Integers, floating-point numbers (real or complex), and complex integers:

```
? factor(239+5*I)
  [-I 1]
  [1 + I 1]
  [117 + 122*I 1]
```

Exact rationals:

```
? 2/3+4/5
  22/15
```

Modular integers:

```
? Mod(3,239)^77
  Mod(128, 239)
```

Vectors and matrices:

```
? v=vector(5,j,j^2)
  [1, 4, 9, 16, 25]
? m=matrix(5,5,r,c,r+c)
  [2 3 4 5 6]
  [3 4 5 6 7]
  [4 5 6 7 8]
  [5 6 7 8 9]
  [6 7 8 9 10]
```

The vector is a row vector, trying to right-multiply it with the matrix fails:

```
? t=m*v
  ***   impossible multiplication t_MAT * t_VEC.
```

The operator ~ transposes vectors (and matrices), we multiply with the column vector:

```
? t=m*v~
%14 = [280, 335, 390, 445, 500]~
```

The result is a column vector, note the tilde at the end of the line.

Vector indices start with one:

```
? t[1]
%15 = 280
```

Symbolic computations

Univariate polynomials:

```
? (1+x)^7
  x^7 + 7*x^6 + 21*x^5 + 35*x^4 + 35*x^3 + 21*x^2 + 7*x + 1
? factor((1+x)^6+1)
  [x^2 + 2*x + 2 1]
  [x^4 + 4*x^3 + 5*x^2 + 2*x + 1 1]
```

Power series:

```
? (1+x+O(x^4))^7
  1 + 7*x + 21*x^2 + 35*x^3 + O(x^4)
? log((1+x+O(x^4))^7)
  7*x - 7/2*x^2 + 7/3*x^3 + O(x^4)
```

Types can be nested, here we compute modulo the polynomial $1 + x + x^7$ with coefficients over GF(2):

```
? t=Mod(1+x, Mod(1,2)*(1+x+x^7))^77
  Mod(Mod(1, 2)*x^3 + Mod(1, 2)*x + Mod(1, 2), Mod(1, 2)*x^7 + Mod(1, 2)*x + Mod(1, 2))
? lift(t)  \\ discard modulo polynomial
  Mod(1, 2)*x^3 + Mod(1, 2)*x + Mod(1, 2)
? lift(lift(t))  \\ discard modulo polynomial, then modulus 2 with coefficient
  x^3 + x + 1
```

Symbolic computations are limited when compared to a computer algebra system: for example, multivariate polynomials cannot (yet) be factored, and there is no symbolic solver for polynomials.

An uninitialized variable evaluates to itself, as a symbol:

```
? hello
  hello
```

To create a symbol, prepend a tick:

```
? w=3
  3
? hello='w  /* the symbol w, not the value of w */
  w
```

Here is a method to create symbols:

```
? sym(k)=eval(Str("A", k))
? t=vector(5, j, sym(j-1))
  [A0, A1, A2, A3, A4]
```

The ingredients are `eval()` and `Str()`:

```
eval(x): evaluation of x, replacing variables by their value.
Str({str}*): concatenates its (string) argument into a single string.
```

Some more trickery to think about:

```
sym(k)=eval(Str("A", k))
t=vector(5, j, sym(j-1));  print("1: t=", t);
{ for (k=1, 5,
    sy = sym(k-1);
    v = 1/k^2;
    /* assign to the symbol that sy evaluates to, the value of v: */
    eval( Str( Str( sy ), "=", Str( v) ) );
 ); }
print("2: t=", t); /* no lazy evaluation with pari/gp */
t=eval(t);   print("3: t=", t);
```

The output of this script is

```
1: t=[A0, A1, A2, A3, A4]
2: t=[A0, A1, A2, A3, A4]
3: t=[1, 1/4, 1/9, 1/16, 1/25]
```

More builtin functions

The following constants and transcendental functions are known by pari:

```
? ?3
  Euler       I           Pi          abs         acos        acosh       agm         arg
  asin        asinh       atan        atanh       bernfrac    bernreal    bernvec     besselh1
  besselh2    besseli     besselj     besseljh    besselk     besseln     cos         cosh
  cotan       dilog       eint1       erfc        eta         exp         gamma       gammah
  hyperu      incgam      incgamc     lngamma     log         polylog     psi         sin
  sinh        sqr         sqrt        sqrtn       tan         tanh        teichmuller theta
  thetanullk  weber       zeta
```

To obtain information about a particular function, use a question mark:

```
? ?sinh
  sinh(x): hyperbolic sine of x.
```

Transcendental functions will also work with complex arguments and symbolically, returning a power series:

```
? sinh(x)
%9 = x + 1/6*x^3 + 1/120*x^5 + 1/5040*x^7 + 1/362880*x^9 \
     + 1/39916800*x^11 + 1/6227020800*x^13 + 1/1307674368000*x^15 + O(x^17)
```

The line break (and the backslash indicating it) was manually entered for layout reasons. The 'precision' (that is default order) of power series can be set by the user:

```
? default(seriesprecision,9);
? sinh(x)
%11 = x + 1/6*x^3 + 1/120*x^5 + 1/5040*x^7 + 1/362880*x^9 + O(x^10)
```

One can also manually give the $O(x^N)$ term:

```
? sinh(x+O(x^23))
%12 = x + 1/6*x^3 + 1/120*x^5 + 1/5040*x^7 + \
   [--snip--] \
  + 1/121645100408832000*x^19 + 1/5109094217170944000*x^21 + O(x^23)
```

Functions operating on matrices are (type `mat`, then hit the tab-key)

```
  matadjoint        matalgtobasis     matbasistoalg     matcompanion
  matdet            matdetint         matdiagonal       mateigen
  matfrobenius      mathess           mathilbert        mathnf
  mathnfmod         mathnfmodid       matid             matimage
  matimagecompl     matindexrank      matintersect      matinverseimage
  matisdiagonal     matker            matkerint         matmuldiagonal
  matmultodiagonal  matpascal         matrank           matrix
  matrixqz          matsize           matsnf            matsolve
  matsolvemod       matsupplement     mattranspose
```

Builtin number theoretical functions are

```
? ?4
  addprimes     bestappr      bezout        bezoutres     bigomega      binomial
  chinese       content       contfrac      contfracpnqn  core          coredisc
  dirdiv        direuler      dirmul        divisors      eulerphi      factor
  factorback    factorcantor  factorff      factorial     factorint     factormod
  ffinit        fibonacci     gcd           hilbert       isfundamental ispower
  isprime       ispseudoprime issquare      issquarefree  kronecker     lcm
  moebius       nextprime     numbpart      numdiv        omega         precprime
  prime         primepi       primes        qfbclassno    qfbcompraw    qfbhclassno
  qfbnucomp     qfbnupow      qfbpowraw     qfbprimeform  qfbred        qfbsolve
  quadclassunit quaddisc      quadgen       quadhilbert   quadpoly      quadray
  quadregulator quadunit      removeprimes  sigma         sqrtint       zncoppersmith
  znlog         znorder       znprimroot    znstar
```

Functions related to polynomials and power series are

```
? ?7
  O                 deriv             eval              factorpadic
  intformal         padicappr         polcoeff          polcyclo
  poldegree         poldisc           poldiscreduced    polhensellift
  polinterpolate    polisirreducible  pollead           pollegendre
  polrecip          polresultant      polroots          polrootsmod
  polrootspadic     polsturm          polsubcyclo       polsylvestermatrix
  polsym            poltchebi         polzagier         serconvol
  serlaplace        serreverse        subst             substpol
  substvec          taylor            thue              thueinit
```

Plenty to explore!

Control structures for programming

Some loop constructs available are

`while(a,seq)`: while a is nonzero evaluate the expression sequence seq. Otherwise 0.

`until(a,seq)`: evaluate the expression sequence seq until a is nonzero.

`for(X=a,b,seq)`: the sequence is evaluated, X going from a up to b.

`forstep(X=a,b,s,seq)`: the sequence is evaluated, X going from a to b in steps of s (can be a vector of steps)

`forprime(X=a,b,seq)`: the sequence is evaluated, X running over the primes between a and b.

`fordiv(n,X,seq)`: the sequence is evaluated, X running over the divisors of n.

The expression `seq` is a list of statements:

```
for ( k=1, 10,  stat1; stat2; stat3; )  /* last semicolon optional */
for ( k=1, 10,  stat1; )
for ( k=1, 10,  ; )  /* zero statements (do nothing, ten times) */
```

(The comments enclosed in `/* */` were manually added.)

The loop-variable is local to the loop:

```
? for(k=1,10, ; )  /* do nothing, ten times */
? k
  k /* not initialized in global scope ==> returned as symbol */
```

A global variable of the same name is not changed:

```
? k=7
  7
? for(k=1,3, print("  k=",k))
  k=1
  k=2
  k=3
? k
  7 /* global variable k not modified */
```

For the sake of clarity, avoid using global and loop-local variables of the same name.

A loop can be aborted with the statement `break()`. The n enclosing loops are aborted by `break(n)`. With `next()`, the next iteration of a loop is started (and the statements until the end of the loop are skipped). With `break(n)` the same is done for the n-th enclosing loop.

And yes, there is an `if` statement:

```
if(a,seq1,seq2): if a is nonzero, seq1 is evaluated, otherwise seq2. seq1 and seq2
    are optional, and if seq2 is omitted, the preceding comma can be omitted also.
```

To have more than one statement in the branches use semicolons between the statements:

```
if ( a==3, /* then */
    b=b+1;
    c=7;
    , /* else */
    b=b-1;
    c=0;
   );
```

Non-interactive usage (scripts)

Usually one will create scripts that are fed into gp (at the command line):

```
gp -q < myscript.gp
```

The option `-q` suppresses the startup message and the history numbers `%N`.

If the script contains just the line

```
exp(2.0)
```

the output would be

```
  7.389056098930650227230427460575007813 1
```

To also see the commands in the output, add a `default(echo,1);` to the top of the file. The output will then be

```
? exp(2.0)
  7.389056098930650227230427460575007813 1
```

You should use comments in your scripts, there are two types of them:

```
\\ a line comment, started with backslashes
/* a block comment
   can stretch over several lines, as in the C-language */
```

Comments are not visible in the output. With the script

```
default(echo, 1);
\\ sum of square numbers:
s=0;  for (k=1, 10, s=s+k*k);  s
```

the output would be

```
? default(echo,1);
? s=0;for(k=1,10,s=s+k*k);s
  385
```

Note that all whitespaces are suppressed in the output.

A command can be broken into several lines if it is inclosed into a pair of braces:

```
{ for (k=1, 10,
       s=s+k*k;
       print(k,":  s=", s);
  ); }
```

This is equivalent to the one-liner

```
 for (k=1, 10,  s=s+k*k;  print(k,":  s=", s);  );
```

User-defined Functions

Now we define a function:

```
powsum(n, p)=
{ /* return the sum 1^p+2^p+3^p+...+n^p */
    local(t);
    t = 0;
    for (k=1, n,
         t = t+k^p);  \\ '^' is the powering operator
    return( t );
}
```

The statement `local(t);` makes sure that no global variable named `t` (if it exists) would be changed by the function. It must be the first statement in the function. The variable `k` in the `for()`-loop is automatically local and should not be listed with the locals. Note that each statement is terminated with a semicolon. The output would be

```
? powsum(n,p)=local(t);t=0;for(k=1,n,t=t+k^p);return(t);
? powsum(10,2)
  385
```

Note how the function definition is changed to a one-liner in the output.

If you have to use global variables, list them at the beginning of your script as follows:

```
global(var1, var2, var3);
```

Any attempt to use the listed names as names of function arguments or local variables in functions will trigger an error.

Arguments are passed by value. There is no mechanism for passing by reference, global variables can be a workaround for this.

Arguments can have defaults, as in

```
powsum(n, p=2)= /* etc */
```

Calling the function as either `powsum(9)` or `powsum(9,)` would compute the number of the first 9 squares. Defaults can appear anywhere in the argument list, as in

```
abcsum(a, b=3, c)= return( a+b+c );
```

So `abcsum(1,,1)` would return 5.

All arguments are implicitly given the default zero, so the sequence of statements

```
foo(a, b, c)= print(a,":",b,":",c);
foo(,,)
foo()
foo
```

will print three times `0:0:0`. This feature is rarely useful and does lead to obscure errors. It will hopefully be removed in future versions of pari/gp.

Bibliography

A

[1] Milton Abramowitz, Irene A. Stegun, (eds.): **Handbook of Mathematical Functions**, National Bureau of Standards, 1964, third printing, (1965). 687, 689, 691, 696

[2] Ramesh C. Agarwal, James W. Cooley: **New algorithms for digital convolution**, IEEE Transactions on Acoustics, Speech, and Signal Processing, vol. ASSP-25, pp.392-410, (October-1977). 543

[3] Simon Joseph Agou, Marc Deléglise, Jean-Louis Nicolas: **Short Polynomial Representations for Square Roots Modulo** p, Designs, Codes and Cryptography, vol.28, pp.33-44, (2003). 773

[4] Manindra Agrawal, Neeraj Kayal, Nitin Saxena: **PRIMES is in P**, Annals of Math., vol.160, no.2, pp.781-793, (September-2004). Online at `http://www.math.princeton.edu/~annals/`. 793

[5] Omran Ahmadi, Alfred Menezes: **On the number of trace-one elements in polynomial bases for** $\mathbb{F}_{2^n}$, Designs, Codes and Cryptography, vol.37, no.3, pp.493-507, (December-2005). Online at `http://citeseer.ist.psu.edu/677930.html`. 885

[6] Omran Ahmadi, Alfred Menezes: **Irreducible polynomials of maximum weight**, Utilitas Mathematica, vol.72, pp.111-123, (2007). Online at `http://www.math.uwaterloo.ca/~ajmeneze/research.html`. 847

[7] W. R. Alford, Andrew Granville, Carl Pomerance: **There are infinitely many Carmichael numbers**, Annals of Mathematics, vol.139, pp.703-722, (1994). 775

[8] Jean-Paul Allouche, Jeffrey Shallit: **The ubiquitous Prouhet-Thue-Morse sequence**, In: C. Ding, T. Helleseth, H. Niederreiter, (eds.), Sequences and Their Applications: Proceedings of SETA'98, pp.1-16, Springer-Verlag, (1999). Online at `http://www.cs.uwaterloo.ca/~shallit/papers.html`. 44, 716

[9] Jean-Paul Allouche, Michael Cosnard: **The Komornik-Loreti constant is transcendental**, Amer. Math. Monthly, vol.107, pp.448-449, (2000). Online at `http://www.lri.fr/~allouche/`. 716

[10] Jean-Paul Allouche, Jeffrey Shallit: **Automatic Sequences**, Cambridge University Press, (2003). 352

[11] Advanced Micro Devices (AMD) Inc.: **AMD Athlon Processor, x86 code optimization guide**, Publication no.22007, Revision K, (February-2002). Online at `http://www.amd.com/`. 19

[12] Advanced Micro Devices (AMD) Inc.: **AMD64 Architecture Programmer's Manual. Volume 3: General-Purpose and System Instructions**, Publication no.24594, Revision 3.14, (September-2007). Online at `http://www.amd.com/`. 27

[13] Advanced Micro Devices (AMD) Inc.: **Software Optimization Guide for AMD64 Processors**, Publication no.25112, Revision 3.06, (September-2005). Online at `http://www.amd.com/`. 7

[14] Ray Andraka: **A survey of CORDIC algorithms for FPGA based computers**, In: Proceedings of the Sixth ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '98), ACM, pp.191-200, (1998). Online at `http://www.andraka.com/papers.htm`. 655

[15] George E. Andrews, Richard Askey, Ranjan Roy: **Special functions**, Cambridge University Press, (1999). 583, 598, 686, 692

[16] Francisco Argüello, Emilio L. Zapata: **Fast Cosine Transform based on the Successive Doubling Method**, Electronics Letters, vol.26, no.19, pp.1616-1618, (September-1990). Online at `http://citeseer.ist.psu.edu/400246.html`. 519

[17] Jörg Arndt, Christoph Haenel: **Pi – Unleashed**, Springer-Verlag, (2000). (translation of: **Pi. Algorithmen, Computer, Arithmetik**, 1998). 603

[18] Jörg Arndt: **Arctan relations for Pi**, (May-2006). Online at `http://www.jjj.de/arctan/arctanpage.html`. 616

[19] Jörg Arndt: **FXT, a library of algorithms**, (1996). Online at `http://www.jjj.de/fxt/`. xi

[20] Jörg Arndt: **hfloat, a library for high precision computations**, (1995). Online at `http://www.jjj.de/hfloat/`. xi, 517

[21] Mikhail J. Atallah, Samuel S. Wagstaff, Jr.: **Watermarking With Quadratic Residues**, In: Proc. of IS-T/SPIE Conf. on Security and Watermarking of Multimedia Contents, SPIE vol.3657, pp.283-288, (1999). Online at `http://homes.cerias.purdue.edu/~ssw/water.html`. 773

B

[22] John C. Baez: **The Octonions**, Bull. Amer. Math. Soc., vol.39, pp.145-205, (2002). Online at `http://math.ucr.edu/home/baez/octonions/`. 812

[23] David H. Bailey: **FFTs in External or Hierarchical Memory**, Journal of Supercomputing, vol.4, no.1, pp.23-35, (March-1990). Online at `http://crd.lbl.gov/~dhbailey/dhbpapers/`. 425

[24] David H. Bailey, P. N. Swarztrauber: **The Fractional Fourier Transform and Applications**, SIAM Review, vol.33, no.3, pp.389-404, (September-1991). Online at `http://crd.lbl.gov/~dhbailey/dhbpapers/`. 442

[25] David H. Bailey, Jonathan M. Borwein, Peter B. Borwein, Simon Plouffe: **The Quest for Pi**, Mathematical Intelligencer, vol.19, no.1, pp.50-57, (January-1997). Online at `http://crd.lbl.gov/~dhbailey/dhbpapers/`. 605

[26] David H. Bailey, Richard E. Crandall: **On the Random Character of Fundamental Constant Expansions**, Experimental Mathematics, vol.10, no.2, pp.175-190, (June-2001). Online at `http://www.expmath.org/`. 622

[27] Robert Baillie, S. Wagstaff, Jr.: **Lucas Pseudoprimes**. Mathematics of Computation, vol.35, no.152, pp.1391-1417, (October-1980). 790

[28] Jean-Claude Bajard, Sylvanus Kla, Jean-Michel Muller: **BKM: A New Hardware Algorithm for Complex Elementary Functions**, IEEE Transactions on Computers, vol.43, no.8, (August-1994). Online at `http://perso.ens-lyon.fr/jean-michel.muller/BKM94.pdf`. 655

[29] Dominique Roelants van Baronaigien: **A Loopless Gray-Code Algorithm for Listing k-ary Trees**, Journal of Algorithms, vol.35, pp.100-107, (2000). 323

[30] Friedrich L. Bauer: **An Infinite Product for Square-Rooting with Cubic Convergence**, The Mathematical Intelligencer, vol.20, pp.12-13, (1998). 675

[31] M. Beeler, R. W. Gosper, R. Schroeppel: **HAKMEM**, MIT AI Memo 239, (29-February-1972). Retyped and converted to html by Henry Baker, (April-1995). Online at `http://home.pipeline.com/~hbaker1/`. 61, 67, 162, 163, 663, 714

[32] Albert H. Beiler: **Recreations in the Theory of Numbers**, Dover Publications, (1964). 806, 810

[33] Hacène Belbachir, Farid Bencherif: **Linear Recurrent Sequences and Powers of a Square Matrix**, INTEGERS, vol.6, (2006). Online at `http://www.integers-ejcnt.org/vol6.html`. 659

[34] Jordan Bell: **Euler and the pentagonal number theorem**, arXiv:math/0510054v2 [math.HO], (17-August-2006). Online at `http://arxiv.org/abs/math.HO/0510054`. 333

[35] A. A. Bennett: **The four term Diophantine arccotangent relation**, Annals of Math., vol.27, no.1, pp.21-24, (September-1925). 614

[36] Lennart Berggren, Jonathan Borwein, Peter Borwein, (eds.): **Pi: A Source Book**, Springer, (1997). 603

[37] E. R. Berlekamp: **Factoring polynomials over finite fields**, Bell System Technical Journal, vol.46, pp.1853-1859, (1967). 851

[38] Bruce C. Berndt, S. Bahrgava, Frank G. Garvan: **Ramanujan's Theories of Elliptic Functions to Alternative Bases**, Trans. Amer. Math. Soc., vol.347, no.11, pp.4163-4244, (1995). Online at `http://www.math.ufl.edu/~frank/publist.html`. 600

[39] Bruce C. Berndt: **Flowers which we cannot yet see growing in Ramanujan's garden of hypergeometric series, elliptic functions and q's**, In: J. Bustoz, M. E. H. Ismail, S. K. Suslov, (eds.): Special Functions 2000: Current Perspective and Future Directions Kluwer, Dordrecht, pp.61-85, (2001). Online at `http://www.math.uiuc.edu/~berndt/publications.html` 600

[40] Pedro Berrizbeitia, T. G. Berry: **Generalized Strong Pseudoprime tests and applications**, Journal of Symbolic Computation, no.11, (1999). 781

[41] Eli Biham, Rafi Chen: **Near-Collisions of SHA-0**, Advances in Cryptology - CRYPTO'04, Lecture Notes in Computer Science, vol.3152, (2004). Online at `http://eprint.iacr.org/2004/146.ps`. 86

[42] Eli Biham, Rafi Chen, Antoine Joux, Patrick Carribault, Christophe Lemuet, William Jalby: **Collisions of SHA-0 and Reduced SHA-1**, Lecture Notes in Computer Science, vol.3494, pp.36-57, (2005). 86

[43] James R. Bitner, Gideon Ehrlich, Edward M. Reingold: **Efficient generation of the binary reflected Gray code and its applications**, Communications of the ACM, vol.19, no.9, pp.517-521, (September-1976). 206, 323

[44] Ian F. Blake, Shuhong Gao, Ronald C. Mullin: **Explicit Factorization of $x^{2^k}+1$ over $\mathbb{F}_p$ with Prime $p \equiv 3 \bmod 4$**, Applicable Algebra in Engineering, Communication and Computation 4, pp.89-94, (1993). Online at `http://www.math.clemson.edu/~sgao/pub.html`. 796

[45] Ian F. Blake, Shuhong Gao, Robert J. Lambert: **Construction and Distribution Problems for Irreducible Trinomials over Finite Fields**, In: D. Gollmann, (ed.): Applications of finite fields, pp.19-32, (1996). Online at `http://www.math.clemson.edu/~sgao/pub.html`. 845

[46] Leo I. Bluestein: **A linear filtering approach to the computation of the discrete Fourier transform**, IEEE Transactions on Audio and Electroacoustics, vol.18, pp.451-455, (December-1970). 441

[47] Antonia W. Bluher: **A Swan-like theorem**, Finite Fields and Their Applications, vol.12, pp.128-138, (28-June-2006). Online at `http://arxiv.org/abs/math/0406538v2`. 885

[48] Arnaud Bodin: **Number of irreducible polynomials in several variables over finite fields**, arXiv:0706.0157v2 [math.AC], (11-June-2007). Online at `http://arxiv.org/abs/0706.0157v2`. 838

[49] Marco Bodrato: **Towards Optimal Toom-Cook Multiplication for Univariate and Multivariate Polynomials in Characteristic 2 and 0**, Lecture Notes in Computer Science, vol.4547, pp.116-133, (2007). Online at `http://bodrato.it/papers/`.825, 885

[50] Marco Bodrato, Alberto Zanoni: **What About Toom-Cook Matrices Optimality?**, Technical Report 605, Centro Vito Volterra, Università di Roma Tor Vergata, (October-2006). Online at `http://bodrato.it/papers/`. 544, 547

[51] Marco Bodrato, Alberto Zanoni: **Integer and Polynomial Multiplication: Towards Optimal Toom-Cook Matrices**, Proceedings of the 2007 international symposium on Symbolic and algebraic computation, pp.17-24, (2007). 545, 547

[52] Miklós Bóna: **Combinatorics of permutations**, Chapmann & Hall/CRC, (2004). 283

[53] J. M. Borwein, P. B. Borwein: **Cubic and higher order algorithms for π**, Canad. Math. Bull., vol.27, pp.436-443, (1984). 607

[54] J. M. Borwein, P. B. Borwein: **More quadratically converging Algorithms for π**, Mathematics of Computation, vol.46, no.173, pp.247-253, (January-1986). 603

[55] J. M. Borwein, P. B. Borwein: **An explicit cubic iteration for π**, BIT Numerical Mathematics, vol.26, no.1, (March-1986). 608

[56] J. M. Borwein, P. B. Borwein: **Pi and the AGM**, Wiley, (1987). 596, 603, 604, 605, 609, 617, 692

[57] J. M. Borwein, P. B. Borwein: **On the Mean Iteration** $(a,b) \leftarrow \left(\frac{a+3b}{4}, \frac{\sqrt{ab}+b}{2}\right)$, Mathematics of Computation, vol.53, no.187, pp.311-326, (July-1989). 598, 609

[58] J. M. Borwein, P. B. Borwein: **A cubic counterpart of Jacobi's Identity and the AGM**, Transactions of the American Mathematical Society, vol.323, no.2, pp.691-701, (February-1991). 598, 607

[59] J. Borwein, P. Borwein, F. Garvan: **Hypergeometric Analogues of the Arithmetic-Geometric Mean Iteration**, Constr. Approx., vol.9, no.4, pp.509-523, (1993). Online at `http://www.math.ufl.edu/~frank/publist.html`. 598, 601

[60] J. M. Borwein, P. B. Borwein, F. G. Garvan: **Some cubic modular identities of Ramanujan**, Transactions of the American Mathematical Society, vol.343, no.1, pp.35-47, (May-1994). Online at `http://citeseer.ist.psu.edu/borwein94some.html`. 604

[61] J. M. Borwein, F. G. Garvan: **Approximations to π via the Dedekind eta function**, Organic mathematics (Burnaby, BC, 1995), pp.89-115, CMS Conf. Proc., vol.20, Amer. Math. Soc., Providence, RI, (1997). Online at `http://www.math.ufl.edu/~frank/publist.html`. 607, 608

[62] Florian Braun, Marcel Waldvogel: **Fast Incremental CRC Updates for IP over ATM Networks**, In: Proceedings of the 2001 IEEE Workshop on High Performance Switching and Routing, (2001). Online at `http://citeseer.ist.psu.edu/braun01fast.html`. 87

[63] Richard P. Brent: **Algorithms for minimization without derivatives**, Prentice-Hall, (1973) (out of print). 582

[64] Richard P. Brent: **Fast multiple-precision evaluation of elementary functions**, Journal of the ACM (JACM), vol.23, no.2, pp.242-251, (April-1976). Online at `http://www.rpbrent.com/`. 603

[65] Richard P. Brent: **On computing factors of cyclotomic polynomials**, Mathematics of Computation, vol.61, no.203, pp.131-149, (July-1993). Online at `http://www.rpbrent.com/`. 791, 885

[66] Richard P. Brent: **On the periods of generalized Fibonacci recurrences**, Mathematics of Computation, vol.63, no.207, pp.389-401, (July-1994). Online at `http://wwwmaths.anu.edu.au/~brent/pub/pubsall.html`. 661

[67] Richard P. Brent: **Computing Aurifeuillian factors**, in: Computational Algebra and Number Theory, Mathematics and its Applications, vol.325, Kluwer Academic Publishers, Boston, pp.201-212, (1995). Online at `http://www.rpbrent.com/`. 791

[68] Richard P. Brent, Alfred J. van der Poorten, Herman J. J. te Riele: **A comparative Study of Algorithms for Computing Continued Fractions of Algebraic Numbers**, In: Lecture Notes in Computer Science, vol.1122, Springer-Verlag, Berlin, pp.35-47, (1996). Online at `http://wwwmaths.anu.edu.au/~brent/pub/pubsall.html`. 704

[69] Richard P. Brent, Samuli Larvala, Paul Zimmermann: **A Fast Algorithm for Testing Irreducibility of Trinomials mod 2 and some new primitive trinomials of degree 3021377**, Mathematics of Computation, vol.72, pp.1443-1452, (2003). Online at `http://www.rpbrent.com/`. 837

[70] Richard P. Brent: **Primality Testing**, Slides of talk to the Oxford University Invariant Society, (25-November-2003). Online at `http://web.comlab.ox.ac.uk/oucl/work/richard.brent/talks.html`. 793

[71] Richard P. Brent, Pierrick Gaudry, Emmanuel Thomé, Paul Zimmermann: **Faster Multiplication in** GF(2)[x], INRIA Technical Report RR-6359, (November-2007). Online at `http://hal.inria.fr/inria-00188261/` and at `http://www.rpbrent.com/`. 826

[72] John Brillhart, Derrick H. Lehmer, John L. Selfridge: **New primality criteria and factorizations of** $2^m \pm 1$, Mathematics of Computation, vol.29, no.130, pp.620-647, (April-1975). 784

[73] J. Brillhart, D. H. Lehmer, J. L. Selfridge, B. Tuckerman, S. S. Wagstaff, Jr.: **Factorizations of** $b^n \pm 1 \quad b = 2, 3, 5, 6, 10, 11$ **up to high powers**, Contemporary Mathematics, Volume 22, second edition, American Mathematical Society, (1988). Online at `http://www.ams.org/`. 786, 793, 837

[74] Bette Bultena, Frank Ruskey: **An Eades-McKay Algorithm for Well-Formed Parenthesis Strings**, Information Processing Letters 68, pp.255-259, (1998). Online at `http://www.cs.uvic.ca/~ruskey/Publications/`. 319

[75] Ronald Joseph Burthe, Jr.: **Further Investigations with the Strong Probable Prime Test**, Mathematics of Computation, vol.65, no.213, pp.373-381, (January-1996). 779

C

[76] David G. Cantor: **On arithmetical algorithms over finite fields**, Journal of Combinatorial Theory, Series A, vol.50, no.2, pp.285-300, (March-1989). 826

[77] K. Cattel, S. Zhang, X. Sun, M. Serra, J. C. Muzio, D. M. Miller: **One-Dimensional Linear Hybrid Cellular Automata: Their Synthesis, Properties, and Applications in VLSI Testing**, tutorial paper, (1996). Online at `http://citeseer.ist.psu.edu/260274.html`. 871

[78] Kevin Cattel, Shujian Zhang: **Minimal Cost One-Dimensional Linear Hybrid Cellular Automata of Degree Through 500**, Journal of Electronic Testing: Theory and Applications, vol.6, pp.255-258, (1995). Online at `http://citeseer.ist.psu.edu/cattell95minimal.html` 867

[79] Kevin Cattell, Frank Ruskey, Joe Sawada, Micaela Serra, C. Robert Miers: **Fast Algorithms to Generate Necklaces, Unlabeled Necklaces, and Irreducible Polynomials over** GF(2), Journal of Algorithms, vol.37, pp.267-282, (2000). Online at `http://www.cis.uoguelph.ca/~sawada/pub.html`. 357, 849

[80] Heng Huat Chan: **On Ramanujan's cubic transformation formula for** ${}_2F_1(\frac{1}{3}, \frac{2}{3}; 1; z)$, Mathematical Proceedings of the Cambridge Philosophical Society, vol.124, pp.193-204, (September-1998). 598

[81] Heng Huat Chan, Kok Seng Chua, Patrick Solé: **Quadratic iterations to** π **associated with elliptic functions to the cubic and septic base**, Trans. Amer. Math. Soc., vol.355, pp.1505-1520, (2003). 598, 608, 609

[82] Howard Cheng, Guillaume Hanrot, Emmanuel Thomé, Eugene Zima, Paul Zimmermann: **Time- and Space-Efficient Evaluation of Some Hypergeometric Constants**, International Symposium on Symbolic and Algebraic Computation, ISSAC'07, pp.85-91, (2007). Online at `http://hal.inria.fr/inria-00177850/`. 636

[83] Busiso P. Chisala: **A Quick Cayley-Hamilton**, The American Mathematical Monthly, vol.105, no.9, pp.842-844, (November-1998). 889

[84] D. V. Chudnovsky, G. V. Chudnovsky: **The computation of classical constants**, Proceedings of the National Academy of Sciences of the United States of America, vol.86, no.21, pp.8178-8182, (1-November-1989). 631

[85] Fan Chun, Persi Diaconis, Ron Graham: **Universal cycles for combinatorial structures**, Discrete Mathematics, vol.11, pp.43-59, (1992). Online at `http://www-stat.stanford.edu/~cgates/PERSI/coauthor.html#Graham`. 863

[86] Jaewook Chung, M. Anwar Hasan: **Asymmetric Squaring Formulae**, Centre for Applied Cryptographic Research, University of Waterloo, Ontario, Canada, (3-August-2006). Online at `http://www.cacr.math.uwaterloo.ca/tech_reports.html`. 544, 545, 546, 550

[87] Frédérick Chyzak, Peter Paule, Otmar Scherzer, Armin Schoisswohl, Burkhard Zimmermann: **The construction of orthonormal wavelets using symbolic methods and a matrix analytical approach for wavelets on the interval**, Experimental Mathematics, vol.10, no.1, pp.67-86, (2001). 538

[88] Mathieu Ciet, Jean-Jacques Quisquater, Francesco Sica: **A Short Note on Irreducible Trinomials in Binary Fields**, (2002). Online at `http://www.dice.ucl.ac.be/crypto/publications/2002/poly.ps`. 844

[89] Douglas W. Clark, Lih-Jyh Weng: **Maximal and Near-Maximal Shift Register Sequences: Efficient Event Counters and Easy Discrete Logarithms**, IEEE Transactions on Computers, vol.43, no.5, pp.560-568, (May-1994). 844

[90] Henri Cohen: **A Course in Computational Algebraic Number Theory**, Springer-Verlag, (1993). Online errata list at `http://www.ufr-mi.u-bordeaux.fr/~cohen/`. 14, 612, 772, 773, 784, 793, 853, 856, 871

[91] Henri Cohen: **On the equation** $a^m + b^n = c^p$, Note, (January-1998). Online at `http://www.math.u-bordeaux.fr/~cohen/fermatgen`. 809

[92] Henri Cohen, Fernando Rodriguez Villegas, Don Zagier: **Convergence acceleration of alternating series**, Experimental Mathematics, vol.9, no.1, pp.3-12, (2000). Online at `http://www.expmath.org/`. 641

[93] Henri Cohen: **Analysis of the sliding window powering method**, Journal of Cryptology, vol.18, no.1, pp.63-76, (January-2005). Preprint titled "Analysis of the flexible window powering method" online at `http://www.math.u-bordeaux.fr/~cohen/`. 557

[94] Stephen A. Cook: **On the minimum computation time of functions**, PhD thesis, Department of Mathematics, Harvard University, (1966). Online at `http://www.cs.toronto.edu/~sacook/`. 543

[95] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein: **Introduction to Algorithms**, MIT Press, second edition, (2001). 123, 157, 542, 572

[96] Richard Crandall, Barry Fagin: **Discrete Weighted Transforms and Large-Integer Arithmetic**, Mathematics of Computation, vol.62, no.205, pp.305-324, (January-1994). 430

[97] Richard E. Crandall, Ernst W. Mayer, Jason S. Papadopoulos: **The twenty-fourth Fermat number is composite**, Mathematicians of Computation, vol.72, no.243, pp.1555-1572, (6-December-2002). 784

[98] Richard Crandall, Carl Pomerance: **Prime Numbers: A Computational Perspective**, second edition, Springer, (2005). 793

[99] R. Creutzburg, M. Tasche: **Parameter Determination for Complex Number-Theoretic Transforms Using Cyclotomic Polynomials**, Mathematics of Computation, vol.52, no.185, pp.189-200, (January-1989). 797

D

[100] Ricardo Dahab, Darrel Hankerson, Fei Hu, Men Long, Julio Lopez, Alfred Menezes: **Software Multiplication using Normal Bases**, IEEE Transactions on Computers, vol.55, pp.974-984, (2006). Online as technical report CACR 2004-12 at `http://www.cacr.math.uwaterloo.ca/`. 891

[101] Ivan Damgård, Peter Landrock, Carl Pomerance: **Average case error estimates for the strong probable prime test**, Mathematics of Computation, vol.61, no.203, pp.177-194, (July-1993). 779

[102] Ivan B. Damgård, Gudmund Skovbjerg Frandsen: **An extended quadratic Frobenius primality test with average and worst case error estimates**, IEEE Transactions on Computers, vol.19, pp.783-793, (February-2003). Online at `http://www.brics.dk/RS/03/9/index.html`. 781

[103] Leonard Eugene Dickson: **History of the Theory of Numbers, vol.I, Divisibility and Primality**, Carnegie Institute of Washington, 1919, unaltered reprint of the AMS, (2002). 680

[104] Leonard Eugene Dickson: **History of the Theory of Numbers, vol.II, Diophantine Analysis**, Carnegie Institute of Washington, 1919, unaltered reprint of the AMS, (2002). 331

[105] Geoffrey Dixon: **Division Algebras, Galois Fields, Quadratic Residues**, arXiv:hep-th/9302113, (23-February-1993). Online at `http://arxiv.org/abs/hep-th/9302113`. 813

[106] P. Duhamel, H. Hollmann: **Split radix FFT algorithm**, Electronics Letters, vol.20, no.1, pp.14-16, (5-January-1984). 412

[107] Pierre Duhamel: **Implementation of "split-radix" FFT algorithms for complex, real and real-symmetric data**, IEEE Transactions on Acoustics, Speech and Signal Processing, vol.34 pp.285-295, (1986). 419

[108] Richard Durstenfeld: **Algorithm 235: random permutation**, Communications of the ACM, vol.7, no.7, p.420, (July-1964). 121

[109] Jacques Dutka: **On Square Roots and Their Representations**, Archive for History of Exact Sciences, vol.36, no.1, pp.21-39, (March-1986). 675

E

[110] Peter Eades, Brendan McKay: **An algorithm for generating subsets of fixed size with a strong minimal change property**, Information Processing Letters, vol.19, p.131-133, (19-October-1984). 183

[111] H.-D. Ebbinghaus, H. Hermes, F. Hirzebruch, M. Koecher, K. Mainzer, J. Neukirch, A. Prestel, R. Remmert: **Zahlen**, second edition, (english translation: **Numbers**), Springer-Verlag, (1988). 61, 812, 813

[112] Karen Egiazarian, Jaako Astola: **Discrete Orthogonal Transforms Based on Fibonacci-type recursions**, IEEE Digital Signal Processing Workshop Proceedings, pp.405-408, (September-1996). Online at `http://citeseer.ist.psu.edu/265285.html`. 498

[113] Gideon Ehrlich: **Loopless Algorithms for Generating Permutations, Combinations, and Other Combinatorial Configurations**, Journal of the ACM, vol.20, no.3, pp.500-513, (July-1973). 254

[114] Doron Zeilberger ('Shalosh B.EKHAD'): **Forty "Strange" Computer-Discovered Hypergeometric Series Evaluations**, (12-October-2004). Online at `http://www.math.rutgers.edu/~zeilberg/pj.html`. 687

F

[115] Steven R. Finch: **Mathematical Constants**, Cambridge University Press, (2003). 333, 692

[116] N. J. Fine: **Binomial Coefficients Modulo a Prime**, The American Mathematical Monthly, vol.54, no.10, pp.589-592, (December-1947). 902

[117] N. J. Fine: **Infinite Products for k-th Roots**, The American Mathematical Monthly, vol.84, no.8. pp.629-630, (October-1977). 675

[118] K. Fong, D. Hankerson, J. López, A. Menezes: **Field inversion and point halving revisited**, Technical Report, CORR 2003-18, Department of Combinatorics and Optimization, University of Waterloo, Canada, (2003). Online at `http://citeseer.ist.psu.edu/fong03field.html`. 878

[119] The Free Software Foundation (FSF): **GCC, the GNU Compiler Collection**, version 4.2.1. Online at `http://gcc.gnu.org/`. 21, 913

[120] The Free Software Foundation (FSF): **Other built-in functions provided by GCC**, section 5.47 of the documentation for GCC version 4.2.1. Online at `http://gcc.gnu.org/onlinedocs/gcc-4.2.1/gcc/Other-Builtins.html`. 21

G

[121] Yves Gallot: **Cyclotomic polynomials and prime numbers**, Note, revised version, (5-January-2001). Online at `http://perso.wanadoo.fr/yves.gallot/papers/cyclotomic.html` 791

[122] Shuhong Gao: **Normal Bases over Finite Fields**, Thesis, University of Waterloo, Ontario, Canada, (1993). Online at `http://www.math.clemson.edu/~sgao/pub.html`. 796

[123] Shuhong Gao, Daniel Panario: **Tests and Constructions of Irreducible polynomials over Finite Fields**, In: F. Cucker, M. Shub, (eds.): Foundations of Computational Mathematics, Springer, pp.346-361, (1997). Online at `http://www.math.clemson.edu/~sgao/pub.html`. 832, 833

[124] Frank Garvan: **Cubic modular identities of Ramanujan, hypergeometric functions and analogues of the arithmetic-geometric mean iteration**, Contemporary Mathematics, vol.166, pp.245-264, (1993). Online at `http://www.math.ufl.edu/~frank/publist.html`. 598

[125] Frank Garvan: **Ramanujan's theories of elliptic functions to alternative bases – a symbolic excursion**, J. Symbolic Computation, vol.20, no.5-6, pp.517-536, (1995). Revised 16-December-2005 edition online at `http://www.math.ufl.edu/~frank/publist.html`. 601

[126] Joachim von zur Gathen, Victor Shoup: **Computing Frobenius maps and factoring polynomials**, Computational Complexity, vol.2, pp.187-224, (1992). Online at `http://www.shoup.net/papers/`. 885

[127] Joachim von zur Gathen, Jürgen Gerhard: **Modern Computer Algebra**, Cambridge University Press, second edition, (2003). List of errata online at `http://www-math.upb.de/mca/`. 678, 793, 856

[128] Willi Geiselmann, Dieter Gollmann: **Self-dual bases in** $\mathrm{GF}(2_{q^n})$, Designs, Codes and Cryptography, vol.3, pp.333-345, (1993). 899

[129] Solomon W. Golomb: **Cyclotomic Polynomials and Factorization Theorems**, The American Mathematical Monthly, vol.85, no.9, pp.734-737, (November-1978). 791

[130] Daniel M. Gordon: **A survey of fast exponentiation methods**, Journal of Algorithms, vol.27, no.1, pp.129-146, (1998). Online at `http://citeseer.ist.psu.edu/gordon97survey.html`. 557

[131] Mark Goresky, Andrew Klapper: **Fibonacci and Galois Representations of Feedback with Carry Shift Registers**, IEEE Transactions on Information Theory, vol.48, no.11, pp.2826-2836, (November-2002). Online at `http://www.math.ias.edu/~goresky/`. 866

[132] Johannes Grabmeier, Alfred Scheerhorn: **Finite Fields in AXIOM**, AXIOM Technical Report no. ATR/5, (1992). Online at `http://citeseer.ist.psu.edu/121340.html`. 893

[133] R. L. Graham, D. E. Knuth, O. Patashnik: **Concrete Mathematics**, second printing, Addison-Wesley, New York, (1988). 282, 320, 680, 686, 689, 711

[134] Torbjörn Granlund, Peter L. Montgomery: **Division by Invariant Integers using Multiplication**, SIGPLAN Notices, vol.29, pp.61-72, (June-1994). Online at `http://swox.com/~tege/`. 5

[135] Torbjörn Granlund: **Instruction latencies and throughput for AMD and Intel x86 processors**, (6-September-2005). Online at `http://swox.com/doc/x86-timing.pdf`. 913

[136] Andrew Granville: **It is easy to determine whether a given integer is prime**, Bulletin (New Series) of the AMS, vol.42, no.1, pp.3-38, (2004). Online at `http://www.dms.umontreal.ca/~andrew/AllPublic.html`. 793

[137] David Gries, Jinyun Xue: **Generating a random cyclic permutation**, BIT Numerical Mathematics, vol.28, no.3, pp.569-572, (September-1988). 122

[138] S. Gudvangen: **Practical Applications of Number Theoretic Transforms**, NORSIG-99, Norwegian Signal Processing Symposium, (1999). Online at `http://www.norsig.no/norsig99/`. 532

[139] Jesús Guillera: **Easy Proofs of Some Borwein Algorithms for π**, arXiv:0803.0991v1 [math.NT], (7-March-2008). Online at `http://arxiv.org/abs/0803.0991v1`. 603

H

[140] Bruno Haible: **CLN, a class library for numbers**, (1996). Online at `http://www.ginac.de/CLN/`. 636

[141] Bruno Haible, Thomas Papanikolaou: **Fast multiprecision evaluation of series of rational numbers**, Technical Report, no. TI-7/97, (18-March-1997). Online at `http://www.informatik.th-darmstadt.de/TI/Veroeffentlichung/TR/Welcome.html`. 635

[142] R. V. L. Hartley: **A More Symmetrical Fourier Analysis Applied to Transmission Problems**, Proceedings of the IRE, vol.30, pp.144-150, (March-1942). 501

[143] B. R. Heap: **Permutations by Interchanges**, Computer Journal, vol.6, pp.293-294, (1963). 249

[144] Gerben J. Hekstra, Ed F. A. Deprettere: **Floating Point Cordic (extended version)**, Technical Report ET/NT 93.15, Delft University of Technology, (8-March-1993). Online at `http://citeseer.ist.psu.edu/hekstra93floating.html`. 655

[145] Nicholas J. Higham: **The Matrix Sign Decomposition and its Relation to the Polar Decomposition**, Linear Algebra and its Applications, 212-213, pp.3-20, (1994). Online at `http://citeseer.ist.psu.edu/higham94matrix.html`. 572

[146] Nicholas J. Higham: **Stable Iterations for the Matrix Square Root**, Numerical Algorithms, vol.15, no.2, pp.227-242, (1997). Online at `http://www.ma.man.ac.uk/~higham/`. 572

[147] Christian W. Hoffmann: **π und das arithmetisch-geometrische Mittel**, Swiss Federal Research Institute WSL, (9-April-2002). Online at `http://www.wsl.ch/staff/christian.hoffmann/pi.pdf`. 608

[148] Alston S. Householder: **Polynomial Iterations to Roots of Algebraic Equations**, Proceedings of the American Mathematical Society, vol.2, no.5, pp.718-719, (October-1951). 580

[149] Alston S. Householder: **The Numerical Treatment of a Single Nonlinear Equation**, McGraw-Hill, (1970). 582, 586, 587, 592

[150] Thomas D. Howell, Jean-Claude Lafon: **The Complexity of the Quaternion Product**, Technical Report TR-75-245, Department of Computer Science, Cornell University, Ithaca, NY, (June-1975). Online at `http://home.pipeline.com/~hbaker1/`. 814

I

[151] F. M. Ives: **Permutation enumeration: four new permutation algorithms**, Communications of the ACM, vol.19, no.2, pp.68-72, (February-1976). 267

J

[152] Gerhard Jaeschke: **On strong pseudoprimes to several bases**, Mathematics of Computation, vol.61, no.204, pp.915-926, (October-1993). 778

[153] T. A. Jenkyns: **Loopless Gray Code Algorithms**, Technical Report CS-95-03, Brock University, Canada, (July-1995). Online at `http://citeseer.ist.psu.edu/134169.html`. 206, 214

[154] S. M. Johnson: **Generation of permutations by adjacent transposition**, Mathematics of Computation, vol.17, no.83, pp.282-285, (July-1963). 254

[155] Dieter Jungnickel, Alfred J. Menezes, Scott A. Vanstone: **On the Number of Self-Dual Bases of** GF(q^m) **Over** GF(q), Proceedings of the American Mathematical Society, vol.109, no.1, pp.23-29, (May-1990). Online at `http://www.math.uwaterloo.ca/~ajmenezes/publications/sdb.pdf`. 898

K

[156] Bahman Kalantari, Jürgen Gerlach: **Newton's Method and Generation of a Determinantal Family of Iteration Functions**, Technical Report DCS-TR 371, Dept. of Computer Science, Rutgers University, New Brunswick, NJ, (1998). Online at `http://citeseer.ist.psu.edu/kalantari98newtons.html`. 586, 587

[157] Bahman Kalantari, Y. Yin: **On Extraneous Fixed-Points of the Basic Family of Iteration Functions**, BIT Numerical Mathematics, vol.43, no.2, pp.453-458, (June-2003). 587

[158] Dan Kalman: **A Singularly Valuable Decomposition: The SVD of a Matrix**, preprint, College Math Journal, vol.27, no.1, (January-1996). Online at `http://www.american.edu/academic.depts/cas/mathstat/People/kalman/pdffiles/index.html`. 574

[159] K. Karamanos: **From symbolic dynamics to a digital approach**, International Journal of Bifurcation and Chaos, vol.11, no.6, pp.1683-1694, (2001). 722

[160] Anatoly A. Karatsuba, Y. Ofman: **Multiplication of multidigit numbers on automata**, Soviet Physics Doklady, vol.7, no.7, pp.595-596, (January-1963). Translated from Doklady Akademii Nauk SSSR, vol.145, no.2, pp.293-294, (July-1962). 542

[161] Richard Kaye: **A Gray Code For Set Partitions**, Information Processing Letters, vol.5, no.6, pp.171-173, (December-1976). Online at `http://www.kaye.to/rick/`. 341

[162] Wilfried Keller: **Fermat factoring status**, (prime factors $k \cdot 2\,n + 1$ of Fermat numbers F_m and complete factoring status). Online at `http://www.prothsearch.net/fermat.html`. 785

[163] Louis V. King: **On the Direct Numerical Calculation of Elliptic Functions and Integrals**, Cambridge University Press, (1924). 603, 623

[164] V. Kislenkov, V. Mitrofanov, E. Zima: **How fast can we compute products?**, in ISSAC: Proceedings of the ACM SIGSAM International Symposium on Symbolic and Algebraic Computation, (1999). Online at `http://citeseer.ist.psu.edu/434789.html`. 557

[165] Andrew Klapper, Mark Goresky: **Feedback Shift Registers, 2-Adic Span and Combiners With Memory**, Journal of Cryptology, vol.10, pp.111-147, (1997). Online at `http://www.math.ias.edu/~goresky/EngPubl.html`. 61, 866

[166] Donald E. Knuth: **Structured programming with `go to` statements**, ACM Computing Surveys, vol.6, no.4, (December-1974). 146

[167] Donald E. Knuth: **The Art of Computer Programming**, third edition, Volume 1: Fundamental Algorithms, Addison-Wesley, (1997). Online errata list at `http://www-cs-staff.stanford.edu/~knuth/`. 648, 650

[168] Donald E. Knuth: **The Art of Computer Programming**, third edition, Volume 2: Seminumerical Algorithms, Addison-Wesley, (1997). Online errata list at `http://www-cs-staff.stanford.edu/~knuth/`. 61, 63, 552, 557, 584, 757, 829, 853

[169] Donald E. Knuth: **The Art of Computer Programming**, second edition, Volume 3: Sorting and Searching, Addison-Wesley, (1997). Online errata list at `http://www-cs-staff.stanford.edu/~knuth/`. 123, 681, 761

[170] Donald E. Knuth: **The Art of Computer Programming**, pre-fascicles for Volume 4. Online at `http://www-cs-staff.stanford.edu/~knuth/`. 83, 182, 189, 224, 225, 254, 274, 276, 283, 863

[171] Wolfram Koepf: **Orthogonal Polynomials and Computer Algebra**, In: R. P. Gilbert et al., (eds.): Recent Developments in Complex Analysis and Computer Algebra, Kluwer, pp.205-234, (1999). Online at `http://www.mathematik.uni-kassel.de/~koepf/Publikationen/index.html`. 600

[172] Kenji Koike, Hironori Shiga: **A three terms Arithmetic-Geometric mean**, Journal of Number Theory, vol.124, pp.123-141, (2007). 598

[173] Richard B. Kreckel: **decimal(γ)** $\approx$ 0.57721566[0 − 9]{1001262760}39288477, web note, (19-January-2008). Online at `http://www.ginac.de/~kreckel/news.html#EulerConstantOneBillionDigits`. 636

[174] Ramanujachary Kumanduri, Cristina Romero: **Number theory with computer applications**, Prentice-Hall, (1998). 703, 773, 776, 801

L

[175] Clement W. H. Lam, Leonard H. Soicher: **Three new combination algorithms with the minimal change property**, Communications of the ACM, vol.25, no.8, pp.555-559, (August-1982). 182

[176] Susan Landau, Neil Immermann: **The Similarities (and Differences) between Polynomials and Integers**, International Conference on Number Theoretic and Algebraic Methods in Computer Science (1993), pp.57-59, version of (8-August-1996). Online at `http://citeseer.ist.psu.edu/landau94similarities.html`. 552

[177] Glen G. Langdon, Jr.: **An algorithm for generating permutations**, Communications of the ACM, vol.10, no.5, pp.298-299, (May-1967). 276

[178] Benoit Leblanc, Evelyne Lutton, Jean-Paul Allouche: **Inverse problems for finite automata: a solution based on Genetic Algorithms**, Lecture Notes in Computer Science vol.1363, pp.157-166, (1998). Online at `http://www.lri.fr/~allouche/bibliorecente.html`. 352

[179] D. H. Lehmer: **On arccotangent relations for π**, American Mathematical Monthly, vol.45, pp.657-664, (1938). 614

[180] D. H. Lehmer: **Interesting series involving the central binomial coefficient**, American Mathematical Monthly, vol.92, no.7, pp.449-457, (1985). 323

[181] Charles E. Leiserson, Harald Prokop, Keith H. Randall: **Using de Bruijn Sequences to Index a 1 in a Computer Word**, MIT Lab for Computer Science, Cambridge, MA, (June-1998). Online at `http://citeseer.ist.psu.edu/leiserson98using.html`. 15

[182] Arjen K. Lenstra: **Integer Factoring**, Designs, Codes and Cryptogaphy, vol.19, pp.101-128, (2000). Online at `http://modular.fas.harvard.edu/edu/Fall2001/124/misc/arjen_lenstra_factoring.pdf`. 793

[183] H. W. Lenstra, Jr., R. J. Schoof: **Primitive Normal Bases for Finite Fields**, Mathematics of Computation, vol.48, no.177, pp.217-231, (January-1987). 893

[184] Qun Ying Liao, Qi Sun: **Normal bases and their dual-bases over finite fields**, Acta Mathematica Sinica, English Series, vol.22, no.3, pp.845-848, (May-2006). 899

[185] Rudolf Lidl, Harald Niederreiter: **Introduction to finite fields and their applications**, Cambridge University Press, revised edition, (1994). 837

[186] J. H. van Lint, R. M. Wilson: **A Course in Combinatorics**, Cambridge University Press, (1992). 372, 863

[187] W. Lipski, Jr.: **More on permutation generation methods**, Computing, vol.23, no.4, pp.357-365, (December-1979). 250, 252

[188] Lisa Lorentzen, Haakon Waadeland: **Continued Fractions and Applications**, North-Holland, (1992). 620, 711

[189] Ya Yan Lu: **Computing the Logarithm of a Symmetric Positive Definite Matrix**, Applied Numerical Mathematics: Transactions of IMACS, vol.26, no.4, pp.483-496, (1998). Online at `http://citeseer.ist.psu.edu/331019.html`. 620

M

[190] Robert S. Maier: **A generalization of Euler's Hypergeometric Transform**, arXiv:math.CA/0302084, (8-April-2003). Online at `http://arxiv.org/abs/math.CA/0302084`. 689, 691

[191] Robert S. Maier: **Algebraic Hypergeometric Transformations of Modular Origin**, arXiv:math.NT/0501425, (25-July-2005). Online at `http://arxiv.org/abs/math.NT/0501425`. 598, 602

[192] Robert S. Maier: **P-symbols, Heun identities, and ${}_3F_2$ identities**, arXiv:0712.4299v2 [math.CA], (30-December-2007). Online at `http://arxiv.org/abs/0712.4299`. 689

[193] Kei Makita, Yasuyuki Nogami, Tatsuo Sugimura: **makita-primedeg-irred-via-aop.pdf**, Electronics and Communications in Japan (Part III: Fundamental Electronic Science), vol.88, no.7, pp.23-32, (2005). 841

[194] H. Malvar: **Fast computation of the discrete cosine transform through fast Hartley transform**, Electronics Letters, vol.22, no.7, pp.352-353, (1986). 519

[195] Henrique S. Malvar: **Fast Computation of the discrete cosine transform and the discrete Hartley transform**, IEEE Transactions on Acoustics, Speech and Signal Processing, ASSP-35, no.10, pp.1484-1485, (October-1987). 519

[196] Keith Matthews: **Solving $Ax^2 - By^2 = N$ in integers, where $A > 0$, $B > 0$ and $D = AB$ is not a perfect square and** $\gcd(A, B) = \gcd(A, N) = 1$, web note, (13-September-2007). Online at `http://www.numbertheory.org/notes.html`. 803

[197] M. D. McIlroy: **A Killer Adversary for Quicksort**, Software Practice and Experience, vol.29, p.1-4, (1999). 125

[198] K. R. McLean: **Divisibility Properties of Binomial Coefficients**, The Mathematical Gazette, vol.58, no.403, pp.17-24, (March-1974). 902

[199] Z. A. Melzak: **Bypasses: a simple approach to complexity**, Wiley, (1983). 334

[200] A. Menezes, P. van Oorschot, S. Vanstone: **Handbook of Applied Cryptography**, CRC Press, (1996). Online at `http://www.cacr.math.uwaterloo.ca/hac/`. 755

[201] Helmut Meyn, Werner Götz: **Self-reciprocal Polynomials Over Finite Fields**, Séminaire Lotharingien de Combinatoire, B21d, 8 pp, (1989). Online at `http://www.emis.de/journals/SLC/`. 842

[202] Peter L. Montgomery: **Modular Multiplication Without Trial Division**, Mathematics of Computation, vol.44, no.170, pp.519-521, (April-1985). 755

[203] Peter L. Montgomery: **A Survey of Modern Integer Factorization Algorithms**, CWI Quarterly, vol.7, no.4, pp.337-365, (1994). Online at `http://citeseer.ist.psu.edu/montgomery94survey.html`. 793

[204] James A. Muir, Douglas R. Stinson: **Minimality and Other Properties of the Width-w Nonadjacent Form**, Technical report CORR 2004-08, Centre for Applied Cryptographic Research (CACR) at the University of Waterloo, Canada, (2004). Online at `http://www.cacr.math.uwaterloo.ca/techreports/2004/tech_reports2004.html`. 65

[205] Jean-Michel Muller: **Elementary Functions: algorithms and implementation**, Birkhäuser, (1997). 647, 655

[206] David R. Musser: **Introspective Sorting and Selection Algorithms**, Software Practice and Experience, vol.8, pp.983-993, (1997). 125

N

[207] H. J. Nussbaumer: **Fast Fourier Transform and Convolution Algorithms**, second edition, Springer-Verlag, (1982). 436

O

[208] C. D. Olds: **Continued Fractions**, The Mathematical Association of America, (1963). 711, 801

[209] Michael Orchard: **Fast bit-reversal algorithms based on index representations in** $\mathrm{GF}(2^n)$, IEEE International Symposium on Circuits and Systems, (1989), vol.3, pp.1903-1906, (1989). 861

P

[210] The PARI Group (PARI): **PARI/GP, version 2.3.3**, (2008). Online at `http://pari.math.u-bordeaux.fr/`. 917

[211] Michael S. Paterson, Larry J. Stockmeyer: **On the number of nonscalar multiplications necessary to evaluate polynomials**, SIAM J. Comput., vol.2, no.1, pp.60-66, (March-1973). 637

[212] W. H. Payne, F. M. Ives: **Combination Generators**, ACM Transactions on Mathematical Software (TOMS), vol.5, no.2, pp.163-172, (June-1979). 182

[213] Carl Pomerance, J. L. Selfridge, S. Wagstaff, Jr.: **The Pseudoprimes to** $25 \cdot 10^9$. Mathematics of Computation, vol.35, no.151, pp.1003-1026, (July-1980). 778, 790

[214] Peter John Potts: **Computable Real Arithmetic Using Linear Fractional Transformations**, Report, Department of Computing, Imperial College of Science, Technology and Medicine, London, (June-1996). Online at `http://citeseer.ist.psu.edu/potts96computable.html`. 687, 691

[215] A. D. Poularikas: **The Transforms and applications handbook**, second edition, CRC Press, (2000). 442

[216] Helmut Prodinger: **On Binary Representations of Integers with Digits -1, 0, +1**, INTEGERS, vol.0, (14-June-2000). Online at `http://www.integers-ejcnt.org/vol0.html`. 65

Q

R

[217] Michael O. Rabin: **Probabilistic algorithms in finite fields**, Technical Report MIT-LCS-TR-213, Massachusetts Institute of Technology, (January-1979). Online at `http://publications.csail.mit.edu/`. 832

[218] Charles M. Rader: **Discrete Fourier Transforms When the Number of Data Samples is Prime**, Proceedings of the IEEE, vol.56, iss.6, pp.1107-1108, (June-1968). 444

[219] Janusz Rajski, Jerzy Tyszer: **Primitive Polynomials Over** GF(2) **of Degree up to 660 with Uniformly Distributed Coefficients**, Journal of Electronic Testing: Theory and Applications, vol.19, pp.645-657, (2003). 848

[220] David Rasmussen, Carla D. Savage, Douglas B. West: **Gray code enumeration of families of integer partitions**, Journal of Combinatorial Theory, Series A, vol.70, no.2, pp.201-229, (1995). Online at `http://www.csc.ncsu.edu/faculty/savage/papers.html`. 332

[221] Phillip A. Regalia, Sanjit K. Mitra: **Kronecker Products, Unitary Matrices and Signal Processing Applications**, SIAM Review, vol.31, no.4, pp.586-613, (December-1989). 453

[222] Hans Riesel: **Prime Numbers and Computer Methods for Factorization**, Birkhäuser, (1985). 788, 793

[223] Frank Ruskey, Carla Savage, Terry MinYih Wang: **Generating Necklaces**, Journal of Algorithms, vol.13, pp.414-430, (1992). 360

[224] Frank Ruskey, Joe Sawada: **An Efficient Algorithm for Generating Necklaces with Fixed Density**, SIAM J. Comput., vol.29, no.2, pp.671-684, (1999). Online at `http://www.cs.uvic.ca/~ruskey/Publications/`. 364

[225] Frank Ruskey, Joe Sawada: **Generating Necklaces and Strings with Forbidden Substrings**, in: Computing and Combinatorics: 6th Annual International Conference, COCOON 2000, Sydney, Australia, Lecture Notes in Computer Science, vol.1858 (2000) pp.330-339, (July-2000). Online at `http://www.cs.uvic.ca/~ruskey/Publications/`. 364

[226] Frank Ruskey, Aaron Williams: **Generating combinations by prefix shifts**, Lecture Notes in Computer Science, 3595, (2005). Extended Abstract for COCOON (2005). Online at `http://www.cs.uvic.ca/~ruskey/Publications/`. 179

S

[227] Eugene Salamin: **Computation of π Using Arithmetic-Geometric Mean**, Mathematics of Computation, vol.30, no.135, pp.565-570, (July-1976). 603

[228] Eugene Salamin: **Application of Quaternions to Computation with Rotations**, Working Paper, Stanford AI Lab, (1979). Edited and TeX-formatted by Henry G. Baker, (1995). Online at `http://home.pipeline.com/~hbaker1/`. 569

[229] Carla D. Savage: **Generating Permutations with k-Differences**, SIAM Journal on Discrete Mathematics, vol.4, no.4, pp.561-573, (November-1990). 276

[230] Carla Savage: **A Survey of Combinatorial Gray Codes**, SIAM Review, vol.39, no.4, pp.605-629, (December-1997). Online at `http://www.csc.ncsu.edu/faculty/savage/papers.html`. 171, 275

[231] Joe Sawada: **Generating Bracelets in Constant Amortized Time**, SIAM J. Comput, vol.31, no.1, pp.259-268, (2001). Online at `http://citeseer.ist.psu.edu/470511.html`. 354

[232] Joe Sawada: **A fast algorithm to generate necklaces with fixed content**, Theoretical Computer Science, vol.301, no.1-3, pp.477-489, (May-2003). 364

[233] Arnold Schönhage: **Schnelle Multiplikation von Polynomen über Körpern der Charakteristik 2**, Acta Informatica, vol.7, no.4, pp.395-398, (December-1977). 826, 885

[234] Ernst Schröder: **On Infinitely Many Algorithms for Solving Equations**, translation by G. W. Stewart of: **Ueber unendlich viele Algorithmen zur Auflösung der Gleichungen**, Mathematische Annalen, (1870). Online at `http://citeseer.ist.psu.edu/93609.html`. 582

[235] Robert Sedgewick: **Permutation Generation Methods**, ACM Computing Surveys (CSUR), vol.9, no.2, pp.137-164, (June-1977). Online at `http://www.princeton.edu/~rblee/ELE572Papers/`. 250

[236] Robert Sedgewick: **Algorithms in C++. Parts 1-4: Fundamentals, Data Structures, Sorting, Searching**, third edition, Addison-Wesley, (1998). 123, 145

[237] Jeffrey Shallit, David Swart: **An Efficient Algorithm for Computing the ith Letter of $\varphi^n(a)$**, In: Proc. 10th Annual SODA conference, pp.768-775, (1999). `http://www.cs.uwaterloo.ca/~shallit/Papers/`. 352

[238] Victor Shoup: **Efficient computation of minimal polynomials in algebraic extensions of finite fields**, In: Proceedings of the 1999 international symposium on Symbolic and algebraic computation, Vancouver, British Columbia, Canada, pp.53-58, 1999 Online at `http://www.shoup.net/papers/`. 884

[239] Victor Shoup: **A Computational Introduction to Number Theory and Algebra**, Cambridge University Press, (2005). Updated version online at `http://www.shoup.net/papers/`. 773

[240] N. J. A. Sloane: **The On-Line Encyclopedia of Integer Sequences**. Online at `http://www.research.att.com/~njas/sequences/`. 11, 44, 45, 51, 62, 63, 65, 72, 73, 74, 75, 76, 82, 94, 138, 139, 249, 282, 284, 285, 286, 300, 302, 304, 305, 307, 309, 310, 312, 320, 322, 332, 334, 336, 338, 344, 346, 347, 357, 361, 362, 371, 372, 390, 680, 682, 699, 713, 714, 715, 716, 717, 718, 719, 720, 722, 726, 729, 731, 732, 740, 741, 742, 743, 744, 745, 764, 767, 773, 775, 778, 786, 791, 803, 807, 812, 834, 838, 842, 844, 847, 848, 866, 874, 893, 898, 900, 902, 903

[241] David M. Smith: **Efficient multiple precision evaluation of elementary functions**, Mathematics of Computation, vol.52, pp.131-134, (1989). 637

[242] Jerome A. Solinas: **Generalized Mersenne Numbers**, Technical report CORR 99-39, University of Waterloo, Canada, (1999). Online at `http://www.cacr.math.uwaterloo.ca/`. 759

[243] Hong-Yeop Song: **Examples and Constructions of Hadamard Matrices**, Dept. of Electical and Electronics Engineering, Yonsei University, Korea, (June-2002). Online at `http://calliope.uwaterloo.ca/~ggong/710T4/Song-lecture.ps`. 372

[244] H. Sorensen, D. Jones, C. Burrus, M. Heideman: **On computing the discrete Hartley transform**, IEEE Trans. on Acoustics, Speech and Signal Processing, vol.33, no.5, pp.1231-1238, (October-1985). 511

[245] Henrik V. Sorensen, Douglas L. Jones, Michael T. Heideman, C. Sidney Burrus: **Real-Valued Fast Fourier Transform Algorithms**, IEEE Transactions on Acoustics, Speech and Signal Processing, vol.35, no.6, pp.849-863, (June-1987). 419

[246] Damien Stehlé, Paul Zimmermann: **A Binary Recursive Gcd Algorithm**, INRIA research report RR-5050, (December-2003). 757

[247] Ralf Stephan: **Divide-and-conquer generating functions. Part I. Elementary sequences**, (2003). Online at `http://arxiv.org/abs/math.CO/0307027`. 727

[248] Carl Størmer: **Sur l'application de la théorie des nombres entiers complexes a la solution en nombres rationnels** $x_1\ x_2\ \ldots x_n\ c_1\ c_2\ \ldots c_n\ k$ **de l'équation:** c_1 **arc tg** $x_1 + c_2$ **arc tg** $x_2 + \ldots + c_n$ **arc tg** $x_n = k\frac{\pi}{4}$, Archiv for Mathematik og Naturvidenskab, B.XIX, Nr.3, (vol.19, no.3), pp.1-96, (1896). 610

[249] Richard G. Swan: **Factorization of polynomials over finite fields**, Pacific J. Math., vol.12, no.3, pp.1099-1106, (1962). 844

[250] Paul N. Swarztrauber: **Bluestein's FFT for Arbitrary N on the Hypercube**, Parallel Computing, vol.17, pp.607-617, (1991). Online at `http://www.cisl.ucar.edu/css/staff/pauls/papers/bluestein/bluestein.html`. 441

T

[251] Tadao Takaoka, Stephen Violich: **Combinatorial Generation by Fusing Loopless Algorithms**, In: Computing: The Australasian Theory Symposium (CATS2006), Hobart, Australia. Conferences in Research and Practice in Information Technology (CRPIT), vol.51. Barry Jay, Joachim Gudmundsson, (eds.), (2006). 319

[252] M. A. Thornton, D. M. Miller, R. Drechsler: **Transformations Amongst the Walsh, Haar, Arithmetic and Reed-Muller Spectral Domains**, International Workshop on Applications of the Reed-Muller Expansion in Circuit Design (RMW), pp.215-225, (2001). Online at `http://engr.smu.edu/~mitch/publications.html`. 476

[253] John Todd: **A problem on arc tangent relations**, Amer. Math. Monthly, vol.56, no.8, pp.517-528, (October-1949). 614

[254] Mikko Tommila: **apfloat, A High Performance Arbitrary Precision Arithmetic Package**, (1996). Online at `http://www.apfloat.org/`. 798

[255] J. F. Traub: **Iterative Methods for the Solution of Equations**, Chelsea, (1964). 582, 583, 591

[256] H. F. Trotter: **Algorithm 115: Perm**, Communications of the ACM, vol.5, no.8, pp.434-435, (August-1962). 254

[257] H. W. Turnbull: **Theory of Equations**, fifth edition, Oliver and Boyd, Edinburgh, (1952). 884

U

V

[258] Vincent Vajnovszki: **Generating a Gray Code for P-Sequences**, Journal of Mathematical Modelling and Algorithms, vol.1, pp.31-41, (2002). 319

[259] Raimundas Vidūnas: **Expressions for values of the gamma function**, arXiv:math.CA/0403510, (30-March-2004). Online at `http://arxiv.org/abs/math/0403510`. 596

[260] San C. Vo: **A Survey of Elliptic Cryptosystems, Part I: Introductory**, NASA Advanced Supercomputing (NAS) Division, (August-2003). Online at `http://www.nas.nasa.gov/News/Techreports/2003/2003.html`. 907

[261] Edward R. Vrscay, William J. Gilbert: **Extraneous fixed points, basin boundaries and chaotic dynamics for Schröder and König rational iteration functions**, Numerische Mathematik, vol.52, no.1, pp.1-16, (January-1987). 586, 587

W

[262] Timothy Walsh: **Generating Gray codes in** $O(1)$ **worst-case time per word**, In: DMTCS 2003, C. S. Calude et al. (eds.), Lecture Notes in Computer Science, vol.2731, pp.73-88, (2003). 171

[263] Xiaoyun Wang, Dengguo Feng, Xuejia Lai, Hongbo Yu: **Collisions for Hash Functions MD4, MD5, HAVAL-128 and RIPEMD**, Cryptology ePrint Archive, Report 2004/199, revised version, (17-August-2004). Online at `http://eprint.iacr.org/2004/199/`. 86

[264] Zhong-De Wang: **New algorithm for the slant transform**, IEEE Transactions Pattern Anal. Mach. Intell. (PAMI), vol.4, no.5, pp.551-555, (September-1982). 473

[265] Henry S. Warren, Jr.: **Hacker's delight**, Addison-Wesley, (2003). 5, 57

[266] André Weimerskirch, Christoph Paar: **Generalizations of the Karatsuba Algorithm for Polynomial Multiplication**, (March-2002). Online at `http://citeseer.ist.psu.edu/571363.html`. 821, 824

[267] André Weimerskirch, Christoph Paar: **Generalizations of the Karatsuba Algorithm for Efficient Implementations**, Technical Report, Ruhr-Universität-Bochum, Germany, corrected version, (2003). Online at `http://www.crypto.ruhr-uni-bochum.de/en_publications.html`. 821

[268] Eric Weisstein: **MathWorld**. Online at `http://mathworld.wolfram.com/`. 337, 687, 742, 902

[269] Mark B. Wells: **Generation of Permutations by Transposition**, Mathematics of Computation, vol.15, no.74, pp.192-195, (April-1961). 252

[270] Mark Weston, Vincent Vajnovszki: **Gray codes for necklaces and Lyndon words of arbitrary base**, Pure Mathematics and Applications (PU.M.A.), vol.17, no.1-2, pp.175-182, (2006). Online at `http://homelinux.capitano.unisi.it/~puma/`. 358

[271] Michael Roby Wetherfield, Hwang Chien-lih: **Lists of Machin-type (inverse integral cotangent) identities for Pi/4**, Online at `http://machination.mysite.wanadoo-members.co.uk/index.html`. 614

[272] Michael Roby Wetherfield: **The enhancement of Machin's formula by Todd's process**, The Mathematical Gazette, vol.80, pp.333-344, (July-1996). 614

[273] E. T. Whittaker, G. N. Watson: **A Course of Modern Analysis**, Cambridge University Press, fourth edition, (1927), reprinted (1990). 618, 687, 689, 692

[274] Mladen Victor Wickerhauser: **Adapted Wavelet Analysis from Theory to Software**, AK Peters, Ltd., Wellesley, Mass., (1994). 536

[275] Herbert S. Wilf: **generatingfunctionology**, second edition, Academic Press, (1992). Online at `http://www.math.upenn.edu/~wilf/DownldGF.html`. 284

[276] Hugh C. Williams: **Éduard Lucas and primality testing**, Wiley, (1989). 793

[277] Mark C. Wilson: **Random and exhaustive generation of permutations and cycles**, arXiv:math/0702753v1 [math.CO], (25-February-2007). Online at `http://arxiv.org/abs/math/0702753`. 122

X

[278] Limin Xiang, Kazuo Ushijima: **On $O(1)$ Time Algorithms for Combinatorial Generation**, The Computer Journal, vol.44, pp.292-302, (2001). 319

Y

Z

[279] S. Zaks: **A new algorithm for generation of permutations**, BIT Numerical Mathematics, vol.24, no.2, pp.196-204, (June-1984). 247

[280] Paweł Zieliński, Krystyna Ziętak: **The Polar Decomposition – Properties, Applications and Algorithms**, Annals of the Polish Mathematical Society, vol.38, (1995). Online at `http://citeseer.ist.psu.edu/zielinski95polar.html`. 570

[281] Neal Zierler: **On a Theorem by Gleason and Marsh**, Proceedings of the American Mathematical Society, vol.9, no.2, pp.236-237, (April-1958). 837

Index

Symbols

−2, representations with radix −2 61
9^{9^9}, computation 440
E, elliptic function 596
K, elliptic function 595
$O(1)$ algorithms 542
$\circledast$, cyclic convolution 427
$\circledast_{\{v\}}$, weighted convolution 436
$\circledast_{lin}$, linear convolution 430
η-product (eta-product) 683
π computation
– AGM vs. binary splitting, 635
– iterative algorithms, 602
$\varphi(n)$, Euler's totient function 763
i^i, computation 623
z^z, series for 699
$\mathcal{W}_v[\,]$, weighted transform 435
2-adic, inverse and square root 59
2D Hilbert curve 56, 162, 351, 734
3D Hilbert curve 351

A

AC (adjacent changes), Gray code 381
`ac_gray_delta()` 382
acceleration of convergence, sumalt algorithm 641
ACF (auto correlation function) 432, 864
acyclic (linear) convolution 430
acyclic (linear) correlation 432
addition, modulo m 753
additive group, with a ring 762
additive inverse, modulo m 756
adjacency matrix of a graph 373
adjacent changes (AC), Gray code 381
adjacent nodes in a graph 373
AGM
– (arithmetic-geometric mean), 593
– 4-th order variant, 594
– and hypergeometric functions, 598
– vs. binary splitting, 635
algebra 810
all-ones polynomial 846, 900
all-ones polynomials, trace vector 885
`all_irredpoly (C++ class)` 850
alternating permutations 284
alternating series, and continued fractions 708
alternating series, sumalt algorithm 641
AND-convolution 480
`apply_permutation()` 112, 120
approximations, initial, for iterations 568
arbitrary length FFT 441
arc (edge) of a digraph 373
arctan relations for π 610
arctan, computation by rectangular scheme 638
argument reduction
– for arctan, 621
– for cos, 625
– for exp, 625
– for log, 621
arithmetic transform 473
arithmetic transform, convolution property 479
arithmetic, modular 753
arithmetic-geometric mean (AGM) 593
array notation 171
array, of bits 160
asm trick, with GCC 4, 521
asymptotics, of an algorithm 541
auto correlation function (ACF) 432, 864
average, of two integers, without overflow 25

B

backtracking 373
base (radix) conversion 636
base field 793
basis functions, Reed-Muller transform 476
Beatty sequence with Φ 743
Bell numbers 139, 336
Bell polynomials 337
Ben-Or test for irreducibility 832
Berlekamp's Q-matrix algorithm 851
Bhaskara equation 801
big endian machine 3
binary

– exponentiation, 555
– finite field, 793, 875
– GCD algorithm, 756
– heap, 157
– powering, 555
– relation, 135
– search, 125
binary polynomial 815
binary splitting
– for rational series, 631
– vs. AGM, 635
– with continued fractions, 707
`binary_necklace (C++ class)` 356
Binet form, of a recurrence 665
binomial coefficient
– and type-2 ONB, 902
– modulo a prime, 902
– number of combinations, 175
bit combinations 66
bit counting 19
bit subsets, via sparse counting 70
bit-array 160
bit-array, fitting in a word 24
bit-block boundaries, determination 13
bit-reversal 35
bit-reversal permutation 93
bit-subset, testing 6
bit-wise
– reversal, 34
– rotation, 28
– zip, 39
`bit_fibgray (C++ class)` 80
`bit_necklace (C++ class)` 31
`bit_rotate_sgn()` 54
`bit_subset (C++ class)` 70
`bitarray (C++ class)` 160
`bitpol_factor()` 855
`bitpol_mult()` 91
`bitpol_normal_q()` 890
`bitpol_refine_factors()` 852
`bitpol_sreduce()` 854
bitrev permutation 93
`BITS_PER_LONG` 3
blocks of bits, counting 21
blocks of bits, creation 12
blocks, swapping via quadruple reversion 100
blue code 49, 360
blue code, fixed points 51
bracelets, as equivalence classes 137
branches, avoiding them 24
bsearch 125
`bsearch()` 125
`bsearch_approx()` 126
`bsearch_ge()` 126
bubble sort 27
builtins, GCC 21
butterfly diagram, for radix-2 transforms 449
byte-wise Gray code and parity 42
`BYTES_PER_LONG` 3

C

C++ class XYZ *see* XYZ (C++ class)
C2RFT *see* real FFT
C2RFT (complex to real FT) 416
canonical sequence 112
carries with mixed radix counting 222
carry, in multiplication 551
CAT, constant amortized time 172
`catalan (C++ class)` 315
Catalan constant 642
Catalan numbers 320, 583
Cayley numbers 811
Cayley-Dickson construction 811
characteristic polynomial
– of a matrix, 889
– of a recurrence relation, 657
– with Fourier transform, 523
characteristic, of a field 876
Chase's sequence, for combinations 189
Chebyshev polynomials
– and Pell's equation, 803
– and products for the square root, 675
– and recurrence for subsequences, 663
– and square root approximants, 674
– as hypergeometric functions, 694
– definition, 667
– fast computation, 672
– with accelerated summation, 642
Chinese Remainder Theorem (CRT) 769
Chinese Remainder Theorem, for convolution 532
chirp z-transform 441
circuit in a graph 373
circulant matrix 894
Clausen's products 692
CLHCA (a class of cellular automata) 871
`clz` (Count Leading Zeros), GCC builtin 21
co-lexicographic order
– (definition), 171
– for combinations, 176
– for compositions, 193
– for permutations, 235
– for subsets of a multiset, 289
– with bit combinations, 66
colex (co-lexicographic) order 171
`comb_rec (C++ class)` 190
`combination_chase (C++ class)` 189
`combination_colex (C++ class)` 177
`combination_emk (C++ class)` 184

`combination_enup (C++ class)` 188
`combination_lex (C++ class)` 176
`combination_pref (C++ class)` 179
`combination_revdoor (C++ class)` 182
combinations, Gray code, with binary words 76
combinations, of k bits 66
combinatorial Gray code 171
companion matrix 658, 889
comparison function, for sorting 129
compiler, smarter than you thought 27
complement, of a permutation 113
complement-shift sequences 379
complementary basis 896
complementing the sequency 48
complete graph 375
complex numbers, construction 793
complex numbers, mult. via 3 real mult. 795
complex numbers, sorting 130
composite modulus 763
compositeness of an integer, test for 775
composition, of permutations 114
`composition_colex (C++ class)` 193
`composition_colex2 (C++ class)` 195
`composition_ex_colex (C++ class)` 195
compositions 193
computation of π, AGM vs. binary splitting 635
concave sequence 140
conditional search, for paths in a graph 380
conditional swap 26
conference matrix 367
conjugates of an element in GF(2^n) 882
connected permutation 285
connection polynomial 857
constant
– Catalan, 642
– CORDIC scaling, 651, 654
– Fibonacci parity, 741
– Gray code, 729
– Komornik-Loreti, 716
– parity number, 714
– Pell, 745
– Pell Gray code, 748
– Pell palindromic, 745
– period-doubling, 722
– rabbit, 740
– revbin, 728
– Roth's, 718
– ruler, 721
– sum of Gray code digits, 731
– sum-of-digits, 727
– Thue, 718
– weighted sum of Gray code digits, 733
– weighted sum-of-digits, 728
constant amortized time (CAT) 172
continued fraction 702
continued fractions, as matrix products 707
convergent, of a continued fraction 702
conversion, float to int 7
conversion, of the radix (base) 636
convex sequence 140
convolution
– acyclic (linear), 430
– and Chinese Remainder Theorem, 532
– and multiplication, 550
– AND-convolution, 480
– by FFT, without revbin permutations, 429
– by FHT, 511
– cyclic, 427
– cyclic, by FHT, 511
– dyadic, 463
– exact, 532
– linear, 430
– mass storage, 439
– negacyclic, 436, 515
– OR-convolution, 479
– property, of the Fourier transform, 428
– right-angle, 436
– skew circular, 436
– weighted, 436
– XOR-convolution, 463
cool-lex, order for combinations 179
Cooley-Tukey FFT algorithm 396
`copy_reverse_0()` 416
copying one bit 8
CORDIC algorithms 650
`coroutine (C++ class)` 165
coroutines 164
correlation 432
`cos_rot()` 518
cosine transform (DCT) 518
cosine, by rectangular scheme 640
cosine, CORDIC algorithm 650
cosine, in a finite field 797
counting bits of a sparse word 20
counting bits of a word 19
counting sort 142
coupled iteration, for the square root 561
CPU instructions, often missed 91
CRC (cyclic redundancy check) 84
`crc32 (C++ class)` 86
`crc64 (C++ class)` 84
cross correlation 432
CRT (Chinese Remainder Theorem) 769
`ctz` (Count Trailing Zeros), GCC builtin 21
cube root extraction 561
cubic convergence 581
cycle in a graph 373
cycle type, of a permutation 283

cycle-leaders, for the Gray code permutation 359
cycle-leaders, for the Gray permutation 106
`cycles (C++ class)` 116
cycles of a permutation and in-place routines 119
cycles, of a permutation 114
cyclic
– convolution, 427
– convolution, by FFT, 429
– correlation, 432
– distance, with binary words, 33
– period, of a binary word, 30
– permutation (definition), 116
– permutation, random, 121
– permutations, and factorial numbers, 243
– permutations, recursive generation, 279
– redundancy check (CRC), 84
– ring, 765
– XOR, 33
`cyclic_perm (C++ class)` 280
cyclotomic polynomials
– (definition), 677
– and primes, 791
– and primitive binary polynomials, 850

D

Daubechies wavelets 538
De Bruijn sequence, to compute bit position 15
De Bruijn graph 377
De Bruijn sequence 207, 862
De Bruijn sequence, as path in a graph 377
De Bruijn sequences, number of 863
`debruijn (C++ class)` 207
decimation in frequency (DIF) 399
decimation in frequency FFT algorithm 399
decimation in time (DIT) 396
decimation in time FFT algorithm 396
delta sequence 434
delta set 171
delta squared process 591
`deque (C++ class)` 154
deque (double-ended queue) 154
derangement 113, 286
derangement order, for permutations 275
DFT (discrete Fourier transform) 393, 396
DIF (decimation in frequency) 399
difference sets, and correlation 434
digraph 373
`digraph (C++ class)` 374
`digraph::sort_edges()` 384
`digraph_paths (C++ class)` 375
`digraph_paths::print_turns()` 384
directed graph 373
discrete cosine transform (DCT) 518
discrete Fourier transform (DFT) 393, 396
discrete sine transform (DST) 519
DIT (decimation in time) 396
division
– algorithm using only multiplication, 559
– CORDIC algorithm, 652
– exact, by $C = 2^k \pm 1$, 60
– exact, with polynomials over GF(2), 820
divisionless iterations for polynomial roots 579
`divisors (C++ class)` 289
Dobinski's formula, for Bell numbers 337
double-ended queue (deque) 154
dragon curve sequence 731
DST (discrete sine transform) 519
`dsth()` 519
dual basis 896
dyadic convolution 463
`dyadic_convolution()` 463

E

E, elliptic function 596
Eades-McKay sequence, for combinations 183
easy case, with combinatorial generation 173
edge of a graph 373
edge sorting, with graph search 384
element of order n 525
elementary functions, as hypergeometric f. 693
elliptic E 596
elliptic K 595
elliptic functions, as hypergeometric functions 700
endian-ness, of a computer 3
endo (Even Numbers DOwn) order 185
endo order, for mixed radix numbers 227
enup (Even Numbers UP) order 185
enup order for combinations 186
enup order, with permutations 270
equivalence classes 135
equivalence relation 135
equivalence relations, number of 139
`equivalence_classes()` 136
Eratosthenes, prime sieve 759
eta-product 683
Euclidean algorithm 756
Euler numbers 285
Euler's identity, for hypergeometric functions 689
Euler's totient function 763
exact convolution 532
exact division 59
exact division, by $C = 2^k \pm 1$ 60
exact division, with polynomials over GF(2) 820
exponent, of a group 762
exponential convergence 581
exponential function
– bit-wise computation, 647
– by rectangular scheme, 640

– computation via $q = \exp(-\pi\, K'/K)$, 623
– iteration for, 623
– of power series, 627
exponentiation
– algorithms, 555
– modulo m, 756
extension field 793, 875
external algorithms 439
extraneous fixed point, of an iteration 587

F

factorial number system 236
factorial, binsplit algorithm for 631
factorization of binary polynomials 851
falling factorial basis 236
fast Fourier transform (FFT) 394
fast Hartley transform (FHT) 501
`fcsr (C++ class)` 864
FCSR (feedback carry shift register) 864
feedback carry shift register (FCSR) 864
Fermat numbers 784
Fermat primes 772
`ffact2cyclic()` 281
`ffs` (Find First Set), GCC builtin 21
FFT
– as polynomial evaluation, 552
– radix-2 DIF, 401
– radix-2 DIT, 399
– radix-4 DIF, 408
– radix-4 DIT, 406
– split-radix algorithm, 411
FFT (fast Fourier transform) 394
FFT caching 557
FFT, for multiplication 550
FFT-primes 526
`fft_complex_convolution()` 430
`fft_dif4l()` 409
`fft_dit4_core_p1()` 407
FHT
– convolution by, 511
– DIF step, 505
– DIF, recursive, 505
– DIT, recursive, 502
– radix-2 DIF, 506
– radix-2 DIT, 503
– radix-2 DIT step, 502
– shift operator, 502
FHT (fast Hartley transform) 501
`fht_dif2()` 506
`fht_dif_core()` 403
`fht_real_complex_fft()` 509
Fibbinary numbers 65, 741
Fibonacci
– k-step sequence, 300
– numbers, 300, 302, 740
– parity, 741
– parity constant, 741
– polynomials, 902
– representation, 741
– setup, of a shift register, 860
– words, 296
– words, Gray code, 79, 296
– words, shifts-order, 210
Fibonacci-Haar transform 496
Fibonacci-Walsh transform 498
FIFO (first-in, first-out), queue 152
filter, for wavelet transforms 534
finite field 793
finite-state machine 162
fixed point, of a function 581
fixed point, of an iteration, extraneous 587
fixed points, of the blue code 51
FKM algorithm 354
FKM algorithm, for binary words 31
four step FFT 424
Fourier shift operator 398
Fourier transform (FT) 393
Fourier transform, convolution property 428
fractional (order) Fourier transform 522
fractional Fourier transform 442
free elements (normal elements) 890
FT (Fourier transform) 393
full path in a graph 373

G

Galois Field 875
Galois setup, of a shift register 860
Gauss' transformation 690
Gaussian normal basis 903
GCC, builtins 21
GCD, computation 756
generator in GF(2^n) 878
generator of a group 762
generator, modulo p 443
generator, program producing programs 520
GF(2^n) (binary finite field) 875
`GF2n (C++ class)` 879, 899
`GF2n::init()` 880, 899
`gf2n_fast_trace()` 877
`gf2n_half_trace()` 888
`gf2n_order()` 878
`gf2n_solve_quadratic()` 886
GNB (Gaussian normal basis) 903
Golay-Rudin-Shapiro sequence 45, 718
Goldschmidt algorithm 574
Gray code
– and radix −2 representations, 62
– binary, reversed, 45

– combinatorial (minimal-change order), 171
– constant, 729
– for bit-subsets, 70
– for combinations, 181
– for combinations of a binary word, 76
– for Fibonacci words, 79, 296
– for Lyndon words, 385
– for mixed radix numbers, 222
– for Pell words, 302, 746
– for sparse signed binary words, 304
– for subsets, with shifts-order, 208
– of a binary word, 40
– permutation, 105
– powers of, 48
– single track, 385
Gray permutation 105
`gray_cycle_leaders (C++ class)` 107
green code 50
grep 147
ground field 793, 875
GRS (Golay-Rudin-Shapiro) sequence 45, 718
`grs_negate()` 454
gslex order, for mixed radix numbers 225

H

`haar()` 484
Haar transform 483
`haar_inplace()` 485
`haar_rev_nn()` 491
Hadamard matrix 365, 812
Hadamard transform 447
half-trace, in GF(2^n) with n odd 888
Halley's formula 583, 586
Hamiltonian cycle 373
Hanoi, towers of, puzzle 723
Hartley shift 502
Hartley transform (HT) 501
hashing, via CRC 84
heap 157
Heap's algorithm for permutations 248
heapsort 142
hexanacci numbers 300, 302
hidden constant, with asymptotics 542
high bits of a word, operations on 16
Hilbert curve
– 3D, 351
– by string substitution, 351
– finite state machine for, 162
– function encoding it, 734
– moves, 56
– turns, 736
homogeneous moves, with combinations 183, 186
homogenous moves, with k-subsets 216
Householder's iteration 586
Householder's method 582
HT (Hartley transform) 501
hyperbolic sine and cosine, by CORDIC 653
hypercomplex numbers 811
hypergeometric function
– (definition), 685
– AGM algorithms, 598
– conversion to continued fraction, 710

I

identical permutation 112
`idsth()` 520
i^i, computation 623
indecomposable permutation 285
index of an element modulo m 762
index of the single set bit 14
index sort 126
infinite products, from series 681
inhomogeneous recurrence 661
initial approximations, for iterations 568
integer partitions 327
integer sequence
– Beatty seq. with Φ, 743
– Carmichael numbers, 775
– Catalan numbers, 320, 583
– Euler function $\varphi(n)$, 764
– Euler numbers, 285
– F-increment RGS, 346
– Feigenbaum symbolic seq., 722
– Fibbinary numbers, 65, 741
– Fibonacci numbers, 300, 302, 740
– fixed points in lex-rev seq., 74
– Golay-Rudin-Shapiro seq., 45, 718
– Gray codes, 729
– GRS (Golay-Rudin-Shapiro) seq., 45, 718
– hexanacci numbers, 300, 302
– hypercomplex multiplication, 812
– indecomposable permutations, 285
– integer partitions, 332
– involutions, 284
– irreducible polynomials, 838
– irreducible self-reciprocal polynomial, 840
– irreducible trinomials, 843
– K-increment RGS, 347
– Kronecker symbols $\left(\frac{-1}{n}\right)$, 732
– Lyndon words, 838
– max-increment RGS, 344
– Mephisto Waltz seq., 717
– Moser – De Bruijn sequence, 62
– necklaces, 361
– non-generous primes, 767
– number of XYZ, *see* number of, XYZ
– numbers both triangular and square, 807
– optimal normal bases, type-1, 900

– optimal normal bases, type-2, 902
– paper-folding seq., 731
– paper-folding seq., signed, 732
– paren words, 82
– partitions into distinct parts, 334
– Pell equation not solvable, 803
– pentanacci numbers, 300, 302
– period-doubling seq., 11, 722
– primes with primitive root 2, 846, 866
– primitive roots of Mersenne primes, 357
– primitive trinomials, 843
– quadratic residues all non-prime, 773
– rabbit seq., 497, 740
– radix −2 representations, 63
– restricted growth strings, 322
– ruler function, 720
– sparse signed binary words, 305
– Stirling numbers of the second kind, 336
– subfactorial numbers, 286
– subset-lex words, 72
– sum of binary digits, 726
– sum of digits of binary Gray code, 731
– swaps with revbin permutation, 94
– tetranacci numbers, 300, 302
– Thue-Morse seq., 44, 450, 713, 812
– tribonacci numbers, 300, 302
– type-1 optimal normal bases, 900
– type-2 optimal normal bases, 902
– values of the Möbius function, 680
– Wieferich primes, 767

integer sequence, by OEIS number
– A000009, 334
– A000010, 764
– A000011, 138
– A000013, 138, 390
– A000029, 138
– A000031, 138, 361
– A000041, 332
– A000045, 300, 302, 304, 307, 309, 740
– A000048, 390, 842
– A000073, 300
– A000078, 300
– A000085, 284
– A000108, 320, 322
– A000110, 139, 336, 346
– A000111, 285
– A000123, 715
– A000129, 744
– A000166, 286
– A000201, 743
– A000213, 302
– A000288, 302
– A000296, 338
– A000322, 302
– A000383, 302
– A000695, 63
– A001037, 362, 838
– A001045, 305, 307
– A001110, 807
– A001122, 847, 866
– A001220, 767
– A001262, 778
– A001333, 304, 744
– A001511, 720
– A001591, 300
– A001592, 300
– A001764, 322
– A002293, 322
– A002294, 322
– A002450, 62
– A002475, 844
– A002812, 786
– A002997, 775
– A003010, 786
– A003188, 729
– A003319, 285
– A003688, 304
– A003714, 65, 741, 742
– A003849, 741
– A004211, 346
– A004212, 346
– A004213, 346
– A005351, 63
– A005352, 63
– A005418, 138, 719
– A005578, 305
– A005614, 740
– A005727, 699
– A005811, 731
– A006130, 307
– A006131, 307
– A006206, 682
– A006498, 310
– A006945, 778
– A007895, 741
– A008275, 282
– A008277, 336
– A008683, 680
– A010060, 44, 714
– A011260, 838
– A014577, 731
– A014578, 718
– A015440, 307
– A015441, 307
– A015442, 307
– A015443, 307
– A015448, 304
– A015449, 304

– A019320, 791
– A020229, 778
– A020985, 45, 718
– A022342, 740
– A027362, 893
– A028859, 309
– A031399, 803
– A032908, 742
– A034947, 732
– A035263, 11, 722, 726
– A036991, 82
– A045687, 94
– A046116, 372
– A046699, 76
– A054639, 902, 903
– A055578, 767
– A055881, 249
– A057460, 844
– A057461, 844
– A057463, 845
– A057474, 845
– A061344, 371
– A064990, 717
– A065428, 773
– A071642, 900, 903
– A072226, 791
– A072276, 778
– A073571, 844
– A073639, 844
– A073726, 844, 874
– A079471, 74
– A079559, 75
– A079972, 312
– A080337, 344
– A080764, 744
– A086347, 309
– A093467, 742
– A095076, 741
– A096393, 357
– A100661, 73
– A101284, 903
– A102467, 834
– A104521, 745
– A106400, 713, 716
– A107220, 848
– A107222, 893
– A107877, 347
– A108918, 72
– A118666, 51
– A118685, 812
– A125145, 309
– A135488, 898
– A135498, 898
– A136250, 903
– A136415, 903
– A136416, 844
– A137310, 903
– A137311, 903
– A137313, 903
– A137314, 903
interior bit blocks, determination 13
interleaving process
– for set partitions, 339
– for Trotter's permutations, 254
interpolation binary search 126
interpolation, linear 126
introsort 125
inverse
– 2-adic, 59
– additive, modulo m, 756
– by exponentiation, 877
– cube root, iteration for, 561
– in GF(Q), 877
– iteration for, 559
– modulo m, by exponentiation, 768
– multiplicative, modulo m, 756
– of a circulant matrix, 894
– permutation, 114
– permutation, in-place computation, 116
– power series over GF(2), 820
– root, iteration for, 565
– square root, iteration for, 560
– XYZ transform, *see* XYZ transform
`inverse_haar()` 485
`inverse_haar_inplace()` 485
inversion formula, Lagrange 584
inversion principle, Möbius 679
invertible modulo m 756
involutions 114, 284
irreducible
– polynomial, 831
– trinomial, 843
`is_cyclic()` 116
`is_quadratic_residue_2ex()` 773
`is_small_prime()` 761
isolated ones or zeros in a word 12
iteration
– and multiple roots, 587
– divisionless, for polynomial roots, 579
– for exp, 623
– for inverse, 559
– for inverse cube root, 561
– for inverse root, 565
– for inverse square root, 560
– for logarithm, 619
– for roots, p-adic, 561
– for the zero of a function, 581
– Goldschmidt, 574

– Householder's, 586
– Schröder's, 582
– synthetic, 713
– to compute π, 602
Ives' algorithm for permutation generation 267

J

Jacobi matrix 538

K

K, elliptic function 595
k-subset 210
Karatsuba multiplication
– for integers, 542
– for polynomials, 821
Komornik-Loreti constant 716
König iteration functions 586
`kronecker()` 772
Kronecker product
– (definition), 451
– of Hadamard matrices, 369
`ksubset_gray (C++ class)` 213
`ksubset_rec (C++ class)` 210
`ksubset_twoclose (C++ class)` 216
Kummer's transformation 693

L

Lagrange inversion formula 584
Lambert series 680, 725, 761
LCM, least common multiple 757
least common multiple (LCM) 757
left-to-right powering 556
Legendre symbol 771
Legendre's relation 597, 702
Lehmer code, of a permutation 237
lex (lexicographic) order 171
lexicographic order
– (definition), 171
– for bit combinations, 67
– for combinations, 176
– for multiset permutations, 290
– for subsets, 201
– for subsets of a binary word, 71
– generalized, for mixed radix numbers, 225
`lfsr (C++ class)` 858
LFSR (linear feedback shift register) 857
LFSR, and Hadamard matrices 365
LHCA (linear hybrid cellular automaton) 866
`lhca2poly()` 868
`lhca_next()` 866
LIFO (last-in, first-out), stack 149
`lin2hilbert()` 162
linear convolution 430
linear correlation 432
linear feedback shift register (LFSR) 857
Linear hybrid cellular automaton (LHCA) 866
linear interpolation 126
linear, function in a finite field 876
Lipski's Gray codes for permutations 250
list recursions, for Gray codes 295
little endian machine 3
localized Hartley transform algorithm 515
localized Walsh transform algorithm 459
logarithm
– bit-wise computation, 645, 648
– computation by rectangular scheme, 638
– computation via AGM, 617
– computation via $\pi/\log(q)$, 617
– curious series for, 622
– iteration using exp, 619
– of power series, 626
loop in a graph 373
loopless algorithm 172
low bits, operations on 9
Lucas test, for primality 788
Lucas-Lehmer test, for Mersenne numbers 786
lucky path, in a graph 384
Lyndon words
– (definition), 354
– and irreducible binary polynomials, 849
– and Mersenne primes, 357
– binary, number of, 361
– number of, 361
– with fixed content, 364
– with fixed density, 362
Lyndon words, as binary word 31
`lyndon_gray (C++ class)` 388

M

m-sequence 365, 862
MAC (modular adjacent changes), Gray code 381
`make_complement_shift_digraph()` 379
`make_complete_digraph()` 374
`make_oddprime_bitarray()` 761
mass storage convolution 439
matrix Fourier algorithm (MFA) 424
matrix square root, applications 569
matrix transposition, and zip permutation 102
matrix transposition, in-place 98
maximal order modulo m 762
`maxorder_element_mod()` 767
mean, arithmetic-geometric 593
median of three elements 124
Mephisto Waltz sequence 717
Mersenne numbers, Lucas-Lehmer test 786
Mersenne primes
– 2^j-th roots, 797
– and Lyndon words, 357

– generalized, 759
– Lucas-Lehmer test, 786
Mersenne-Walsh transform 498
MFA (matrix Fourier algorithm) 424
minimal polynomial, in GF(2^n) 882
minimal-change order *see* Gray code
minimum, among bit-wise rotations 30
`minweight_lhca_rule()` 867
missing, CPU instructions 91
mixed radix numbers 219
`mixedradix_endo (C++ class)` 227
`mixedradix_endo_gray (C++ class)` 230
`mixedradix_gray (C++ class)` 222
`mixedradix_gslex (C++ class)` 225
`mixedradix_gslex_alt (C++ class)` 227
`mixedradix_lex (C++ class)` 219
`mixedradix_modular_gray (C++ class)` 225
Möbius function 679
Möbius inversion principle 679
`mod (C++ class)` 527, 798
mod m FFTs 525
modular adjacent changes (MAC), Gray code 381
modular arithmetic 753
modular multiplication 754
modular reduction, with structured primes 759
modular square root 773
modulo, as equivalence classes 136
modulus
– composite, 763
– prime, 762
– prime, with NTTs, 526
moment conditions, for wavelet filters 536
monotone sequence 139
Moser – De Bruijn sequence 62
moves, of the Hilbert curve 56
`mpartition (C++ class)` 331
`mset_lex (C++ class)` 292
`mset_lex_rec (C++ class)` 291
multi-dimensional Walsh transform 450
multi-point iteration 591
multigrades 804
multigraph 373
multinomial coefficient 290, 364
multiple roots, iterations for 587
multiplication
– by FFT, 550
– carry, 551
– integer vs. float, 6
– is convolution, 550
– Karatsuba, 542
– modulo m, 754
– of complex numbers via 3 real mult., 795
– of hypercomplex numbers, 811
– of octonions, 813
– of polynomials, 431
– of quaternions, 813
– sum-of-digits test, 554
multiplication matrix, for normal bases 890
multiplication table, of an algebra 810
multiplicative function 679, 764
multiplicative group, with a ring 762
multiplicative inverse, modulo m 756
multiset 289

N

N-polynomial (normal polynomial) 890
NAF (nonadjacent form) 64
NAF, Gray code 304
`necklace (C++ class)` 207, 355
`necklace2bitpol (C++ class)` 849
necklaces
– as equivalence classes, 137
– binary, 30
– binary, number of, 361
– definition, 353
– with fixed content, 364
– with fixed density, 362
necklaces, as binary words 31
negacyclic convolution 436, 515
neighbors of a node in a graph 373
Newton's formula 884
Newton's iteration, for vector-valued functions 538
node (vertex) of a graph 373
non-generous primes 767
nonadjacent form (NAF) 64
nonadjacent form (NAF), Gray code 304
noncyclic ring 765
normal bases, for GF(2^n) 890
normal basis, optimal 900
normal polynomial 890
`normal_mult()` 890
`normal_solve_reduced_quadratic()` 891
NTT
– (number theoretic transforms), 525
– radix-2 DIF, 528
– radix-2 DIT, 527
– radix-4, 530
number of
– alternating permutation, 284
– alternating permutations, 285
– aperiodic necklaces, 361
– binary necklaces, 361
– binary partitions of even numbers, 715
– binary reversible strings, 138
– binary words at most r successive ones, 299
– bracelets, 138
– carries, 222
– cycles in De Bruijn graph, 379

– De Bruijn sequences, 863
– derangements, 286
– equivalence relations, 139
– F-increment RGS, 346
– fixed density Lyndon words, 362
– fixed density necklaces, 362
– generators modulo n, 767
– increment-i RGS, 322
– indecomposable permutations, 285
– integer partitions, 332
– integers coprime to n, 763
– invertible circulant matrices, 894
– involutions, 284
– irreducible polynomials, 838
– irreducible SRPs, 842
– K-increment RGS, 347
– Lyndon words, 361, 838
– m-sequences, 862
– max-increment RGS, 344
– necklaces, 138, 361
– normal polynomials, 892, 895
– ones in binary Gray code, 731
– parenthesis pairs, 320
– partitions into distinct parts, 334
– permutations of a multiset, 290
– permutations with m cycles, 282
– primitive normal polynomials, 893
– primitive polynomials, 838
– self-dual normal bases, 898
– shift register sequences, 862
– sparse signed binary words, 305
– strings with fixed content, 364
– swaps with revbin permutation, 94
– units in GF(Q), 878
– units modulo m, 764
– unlabeled bracelets, 138
– unlabeled necklaces, 138
number theoretic transforms (NTT) 525

O

$O(1)$ algorithm 172
octonions 811
ONB (optimal normal basis) 900
one-point iteration 581
optimal normal basis (ONB) 900
optimization, with combinatorial generation 172
OR-convolution 479
order
– of a polynomial, 834
– of an element modulo m, 761
– of an iteration, 581
out of core algorithms 439

P

p-adic roots, iterations for 561
Padé approximants
– for arctan, 620
– for exp, 624
– for the logarithm, 619
paper-folding sequence 731
`paren (C++ class)` 313
`paren_gray (C++ class)` 319
parentheses, and binary words 81
parity
– number, 714, 726
– of a binary word, 42
– of a permutation, 116
`parity` (parity of a word), GCC builtin 21
Parseval's equation 394
partition
– of a set, 335
– of an integer, 327
`partition (C++ class)` 329
`partition_rec (C++ class)` 327
partitioning, for quicksort 124
Pascal's triangle 176
path in a graph 373
`pcrc64 (C++ class)` 87
Pell
– constant, 745
– equation, 801
– Gray code constant, 748
– palindromic constant, 745
– ruler function, 746
Pell words, Gray code 302, 746
pentagonal number theorem 333, 651
pentanacci numbers 300, 302
pentanomial 845
Pepin's test, for Fermat numbers 784
period of a polynomial 834
period-doubling constant 722
period-doubling sequence 11, 722
`perm_colex (C++ class)` 235
`perm_derange (C++ class)` 275
`perm_gray_ffact (C++ class)` 260
`perm_gray_ffact2 (C++ class)` 258
`perm_gray_lipski (C++ class)` 250
`perm_gray_rfact (C++ class)` 260
`perm_gray_rot1 (C++ class)` 264
`perm_gray_wells (C++ class)` 252
`perm_heap (C++ class)` 249
`perm_heap2 (C++ class)` 249
`perm_heap2_swaps (C++ class)` 250
`perm_ives (C++ class)` 267
`perm_lex (C++ class)` 233
`perm_mv0 (C++ class)` 265
`perm_rec (C++ class)` 278

`perm_restrpref (C++ class)` 283
`perm_rev (C++ class)` 245
`perm_rev2 (C++ class)` 247
`perm_rot (C++ class)` 276
`perm_st (C++ class)` 269
`perm_st_gray (C++ class)` 272
`perm_star (C++ class)` 274
`perm_star_swaps (C++ class)` 275
`perm_trotter (C++ class)` 254
`perm_trotter_lg (C++ class)` 257
permutation
– alternating, 284
– as path in the complete graph, 377
– composition, 114
– connected, 285
– cycle type, 283
– cycles, 114
– cyclic, random, 121
– derangement, 286
– indecomposable, 285
– inverse of, 114
– involution, 114, 284
– of a multiset, 290
– random, 121
– with m cycles, number of, 282
Pfaff's reflection law 689
phi function, number thoeretic 763
π, computation 602
pitfall, shifts in C 5
pitfall, two's complement 5
Pocklington-Lehmer test, for primality 783
pointer sort 128
pointer, size of 3
polar decomposition, of a matrix 570
polynomial
– binary, weight, 842
– irreducible, 831
– multiplication, 431
– multiplication, splitting schemes, 821
– primitive, 835
– roots, divisionless iterations for, 579
polynomial, binary 815
`popcount` (bit-count), GCC builtin 21
power series
– computation of exponential function, 627
– computation of logarithm, 626
– reversion, 585
powering
– algorithms, 555
– modulo m, 756
– of permutations, 118
– of the binary Gray code, 48
Pratt's certificate of primality 781
prefix shifts, order for combinations 179
`prev_lexrev()` 463
prime length FFT, Rader's algorithm 443
primes
– and cyclotomic polynomials, 791
– as modulus, 762
– as modulus, with NTTs, 526
– non-generous, 767
– sieve of Eratosthenes, 759
– structured, 759
– Wieferich, 767
– with primitive root 2, 866
primitive
– n-th root, modulo m, 525
– r-th root of unity, 761
– elements of a group, 762
– pentanomial, 845
– polynomial, 835
– root, 443, 762
– root in GF(2^n), 878
– root of Mersenne primes, 357
– trinomial, 843, 874
`print_cycles()` 115
priority queue 157
`priority_queue (C++ class)` 158
product form
– for a-th root, 577
– for continued fractions, 707
– for elliptic K, 596
– for power series of exp, 627
– for square root, 675
products of k out of n factors 178
products, infinite, from series 681
Proth's theorem 784
Prouhet-Thue-Morse constant 714
pseudo graph 373
pseudo-inverse, of a matrix 572
pseudo-irreducible (PI) 833
pseudoprime 775
pseudoprime, strong (SPP) 775
Pythagorean triples 805

Q

Q-matrix 851
quadratic convergence 581
quadratic residue (square) modulo p 771
quadratic residues, and Hadamard matrices 367
quadruple reversion trick 100
quantization 134
`quantize()` 134
quartic convergence 581
quaternions 811
`queue (C++ class)` 153
queue (FIFO) 152
quicksort 124

R

R2CFT *see* real FFT
R2CFT (real to complex FT) 416
rabbit constant 740
rabbit sequence 497, 740
Rabin's test for irreducibility 832
Rabin-Miller test, for compositeness 777
Rader's algorithm, for prime length FFT 443
radix −2 representations 61
radix (base) conversion 636
radix permutation 97
radix sort 144
`radix_sort()` 144
random permutation 121
random selection 121
ranking, with combinatorial objects 172
rational, square root iterations 562
re-orthogonalization, of a matrix 569
real FFT
– by FHT, 509
– split-radix algorithm, 420
– with wrap routines, 418
reciprocal polynomial 838
rectangular scheme
– for arctan and log, 637
– for exp, sin, and cos, 639
recurrence
– (definition), 657
– inhomogeneous, 661
– relation, 657
– relation, for subsequences, 663
red code 50
reduction
– modular, with structured primes, 759
– modulo $x^2 + x + 1$ etc., 795
Reed-Muller transform
– (definition), 476
– and necklaces, 359
– convolution property, 480
relation, binary 135
relex order 171
representations, radix −2 61
representatives, with equivalence classes 136
restricted growth strings
– (definition), 315
– for k-ary trees, 321
– for parenthesis strings, 315
– for set partitions, 340
revbin
– (bit-wise reversal), 34
– constant, 728
– pairs, via shift registers, 861
– permutation, 93
– permutation, and convolution by FFT, 429
`revbin_permute()` 95
`revbin_permute0()` 96
reversal, of a permutation 114
`reverse_0()` 416
reversed Gray code 45
reversed Gray code permutation 109
reversed zip permutation 103
reversing the bits of a word 34, 35
reversion of power series 583
reversion, of series, for Schröder's formula 585
RGS (restricted growth string) 315
`rgs_binomial (C++ class)` 321
`rgs_fincr (C++ class)` 346
`rgs_maxincr (C++ class)` 342
right-angle convolution 436
right-to-left powering 555
ring buffer 151
ring, cyclic 765
`ringbuffer (C++ class)` 151
rising factorial basis 236
root
– p-adic, iterations for, 561
– extraction, 565
– inverse, iteration for, 565
– of a polynomial, divisionless iterations, 579
– primitive, 762
– primitive, in GF(2^n), 878
– primitive, modulo m, 525
– primitive, of Mersenne primes, 357
`rotate_left()` 99
`rotate_sgn()` 478
rotation, bit-wise 28
rotation, by triple reversion 99
row-column algorithm 424
ruler constant 721
ruler function 206, 720
`ruler_func (C++ class)` 206, 287

S

Sande-Tukey FFT algorithm 399
scalar multiplication 875
Schröder's formula 582
search, binary 125
searching, with unsorted arrays 145
secant method 581
sedenions 811
selection sort 123
self correlation 432
self-dual (basis over GF(2^n)) 896
self-reciprocal polynomial 840
semi-continuous Fourier transform 395
sentinel element 173
sequence *see* integer sequence
sequency 466

sequency, of a binary word 46
series reversion 583
series reversion, for Schröder's formula 585
set partition 335
`set_partition (C++ class)` 339
`set_partition_rgs (C++ class)` 341
`setup_q_matrix()` 851
shift operator, for Fourier transform 398
shift operator, for Hartley transform 502
shift register sequence (SRS) 857
shift-and-add algorithms 645
shifts in C, pitfall 5
shifts, and division 4
shifts-order
– for bit combinations, 68
– for subsets, 208
sieve of Eratosthenes 759
sign decomposition, of a matrix 571
sign of a permutation 116
sign of the Fourier transform 394
signed binary representation 64
signed binary words, sparse, Gray code 304
simple continued fraction 703
simple path in a graph 373
sine transform (DST) 519
sine, CORDIC algorithm 650
sine, in a finite field 797
single track
– binary Gray code, 385
– order for permutations, 269
– order for subsets, 207
singular value decomposition (SVD) 570
skew circular convolution 436
slant transform 472
slant transform, sequency ordered 472
`slow_convolution()` 427
`slow_ft()` 394
smart, your compiler 27
sorting, edges in a graph 384
sorting, using a heap 142
space-filling Hilbert curve 351
sparse counting, and bit subsets 70
sparse signed binary representation 64
sparse signed binary words, Gray code 304
sparse words, bit counting 20
split-radix FFT algorithm 411
splitting schemes for multiplication
– for integers, 542
– for polynomials over GF(2), 821
splitting, binary, for rational series 631
SPP (strong pseudoprime) 775
`sqrt_modf()` 775
`sqrt_modp()` 773
`sqrt_modpp()` 774
square modulo p 771
square of a permutation 117
square root
– 2-adic, 60
– in GF(2^n), 878
– iteration for, 560
– modulo p, 773
– of a matrix, applications, 569
squarefree factorization, with polynomials 856
SRS (shift register sequence) 857
stable sort 143
`stack (C++ class)` 149
stack (LIFO) 149
stack, for coroutine emulation 165
star-transposition order, for permutations 274
state engine, for coroutine emulation 165
state-engine 162
Stirling numbers
– of the first kind (cycle numbers), 282, 695
– of the second kind (set numbers), 336
strictly convex sequence 141
strictly monotone sequence 140
string substitution engine 349
`string_subst (C++ class)` 350
strings with fixed content 364
strong minimal-change order 171, 254, 319
strong minimal-change order for combinations 183
strong pseudo-irreducible (SPI) 833
strong pseudoprime (SPP) 775
structured primes 759
subdegree of a polynomial 845
subfactorial numbers 286
subsequences, recurrence relations for 663
subset of bit-set, testing 6
`subset_debruijn (C++ class)` 207
`subset_deltalex (C++ class)` 202
`subset_gray (C++ class)` 205
`subset_gray_delta (C++ class)` 203
`subset_lex (C++ class)` 202
subsets
– of k bits (combinations), 66
– of a binary word, 70, 71
– of a multiset, 289
subtraction, modulo m 753
sum of Gray code digits constant 731
sum of two squares 799
sum-of-digits constant 727
sum-of-digits test, with multiplication 554
sumalt algorithm 641
`sumdiff()` 420, 455
super-linear iteration 581
SVD (singular value decomposition) 570
Swan's theorem 844
swap without temporary 7

swap, conditional 26
swapping blocks via quadruple reversion 100
swapping two bits 9
symmetries, of revbin permutation 94
synthetic iterations 713

T

`tcrc64 (C++ class)` 86
tensor product 451
tetranacci numbers 300, 302
Thue constant 718
Thue-Morse sequence 44, 450, 713, 812
`thue_morse (C++ class)` 44
TMFA (transposed matrix Fourier algorithm) 425
toggle between two values 6
Toom-Cook algorithm 543
Toom-Cook algorithm for binary polynomials 825
totient function 763
towers of Hanoi 723
trace
– of a polynomial, 889
– of an element in GF(2^n), 877
– vector, fast computation, 884
– vector, in finite field, 877
trace-orthonormal basis 896
transformations, for elliptic K and E 700
transformations, of hypergeometric functions 688
transforms, on binary words 49
transition count, for a Gray code 385
`transpose()` 98
transposed matrix Fourier algorithm (TMFA) 425
transposed Reed-Muller transform 477
transposition of a matrix, and zip permutation 102
transposition of a matrix, in-place 98
transposition, with permutation 116
`tree_gray (C++ class)` 323
triangular square numbers 807
tribonacci numbers 300, 302
trigonometric recursion 402
trinomial
– primitive, 843, 874
– trace vector, 885
triple reversion trick 99
Trotter's algorithm for permutations 254
two's complement, pitfall 5
two-close order for k-subsets 216
two-close order for combinations 186
type, of a set partition 337
type-1 optimal normal basis 900
type-2 optimal normal basis 901
type-t Gaussian normal basis 903

U

unique 132
units (invertible elements) 756
unlabeled bracelets 138
unranking, with combinatorial objects 172
unsorted arrays, searching 145
unzip permutation 102
`unzip_rev()` 104

V

vertex, of a graph 373
vertical addition 22

W

Walsh transform 447
Walsh transform, multi-dimensional 450
`walsh_pal_basefunc()` 466
`walsh_q1()` 470
`walsh_q2()` 470
`walsh_wak_basefunc()` 449
`walsh_wak_dif2()` 449, 476
`walsh_wak_dit2()` 448, 477
`walsh_wal_basefunc()` 467
`walsh_wal_rev()` 468
`walsh_wal_rev_basefunc()` 469
wavelet conditions 534
wavelet filter 534
wavelet transform 533
`wavelet_filter (C++ class)` 535, 537
weight, of binary polynomial 842
weighted convolution 436
weighted sum of Gray code digits constant 733
weighted sum-of-digits constant 728
weighted transform 435
Wells' Gray code for permutations 252
Whipple's identity 690
Wieferich primes 767

X

XOR permutation 104
XOR, cyclic 33
XOR-convolution 463
`xor_permute()` 104
`xrevbin()` 55
x^x, series for 699

Y

yellow code 49, 359

Z

Z-order 57
z-transform 441
Zeckendorf representation 741
zero bytes, finding 58
zero divisors, of an algebra 810
zero padding, for linear convolution 430
zero-one transitions, determination 13

`zip()` 101
zip permutation 101
zip, bit-wise 39
`zip_rev()` 103
z^z, series for 699