Power Generation Technologies

Second edition

Power Generation Technologies

Second edition

by

Paul Breeze

ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SYDNEY • TOKYO
Newnes is an imprint of Elsevier

Newnes

Newnes is an imprint of Elsevier
The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, UK
225 Wyman Street, Waltham, MA 02451, USA

Second Edition 2014

First Edition 2005

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data
A catalog record for this book is available from the Library of Congress

Printed and bound in United Kingdom

14 15 16 17 18 10 9 8 7 6 5 4 3 2 1

ISBN: 978-0-08-098330-1

Chapter 1

An Introduction to Electricity Generation

Electricity is at the root of everything that we think of as modern. In a practical sense it defines modernity. All of those adjuncts to living in an advanced society that began to appear from the end of the 19th century—electric lighting then electric motors, radio, television, home appliances, and, in the last part of the 20th century, the myriad of electronic devices that have been spawned by the development of the transistor including computers and portable telephones—rely exclusively on electricity for their operation. Their widespread use would not be possible without electricity and the complex electricity supply system that has evolved to deliver it.

Not only is electricity one of the foundations of a modern developed society, electricity is also capable of nourishing the advancement of a society. Something as simple as the availability of electric lighting can lead to enormous benefits in terms of levels of education and quality of life. In consequence, electricity supply is a key element of international development aid. Meanwhile the citizens of many less-developed nations yearn for an adequate electricity supply and all the benefits that it can bring. Ironically, most of the citizens of the world's advanced societies take it for granted.

The industry that supplies electricity and maintains the network that allows it to be delivered to virtually any location on the planet makes up what is probably the largest single industrial endeavor in the world. At the same time, the supply of electricity is a complex operation. Electricity is not a physical commodity like steel or maize even though it is often bought and sold as if it were such a commodity. Electricity is an ephemeral energy source that must be consumed immediately after it is produced. This means that any power station that is producing electrical power must have a customer ready to use it. This careful balancing act is carried out across a network of electricity supply lines controlled by network operators whose primary job is to ensure that the balance between demand and supply is maintained at all times.

Electricity supply is also a security issue. While people untouched by modernity can still live their lives without electricity, a modern industrial nation deprived of its electricity supply is like a great ocean liner without it engines. It becomes helpless. Consequently, governments must ensure that their people and their industries are kept supplied, and national electricity supply strategies will often have security of supply as one of their main considerations.

This book is primarily about the ways of generating electricity. It does not cover in depth the means of transporting electricity and delivering it to those who wish to use it. Nor does it treat, except obliquely, the political issues that attach themselves to electricity supply. What it does attempt is to provide an explanation of all the myriad ways that humans have devised to produce this most elusive of energy forms.

The book is divided into chapters each devoted to one type of electricity generation. The explanations provided are thorough and technical where necessary, but do not resort to overly technical language where it can be avoided. Readers who are seeking a full analysis of the thermodynamics of the heat engine or the differential equations for solving the problem of turbine flow, will need to look elsewhere, but those who seek a thorough understanding of electricity generation will find it here.

The aim of this book is to provide a description of every type of power generation. Even so, there will be occasional lacunas; there is no description of magnetohydrodynamic power generation, for example, although even this obscure phenomenon does earn a brief mention in Chapter 14 on marine power generation. That aside, all practical and some still experimental means of producing electricity are included.

HISTORY OF ELECTRICITY GENERATION

The roots of the modern electricity-generating industry are found in the early and middle years of the 19th century and in the work of men such as André Ampère, Michael Faraday, Benjamin Franklin, and Alessandro Volta. It was during this period that scientists began to forge an understanding of the nature of electrical charge and magnetic fields. The chemical battery that converted chemical energy into electricity had also been discovered and permitted the properties of a flowing electrical charge (an electric current) to be explored. This also allowed the development of the telegraph, the first electrical means of communication. It was Faraday who was able to establish the relationship between electric currents and magnetism, a relationship that makes it possible to generate electricity with moving machinery rather than taking it exclusively from chemical batteries. His discoveries opened the way to the use of rotating engines as a source of electrical power.

The widening understanding of electricity coincided with the development of the steam engine as well as the widespread use of gas for fuel and lighting. Lighting, in particular, caught the public's imagination and one of the first major uses for electricity was as a source of light. In the United States, Thomas Edison developed the carbon filament that produced light from an electric current. Similar work was carried out in the United Kingdom by Sir Joseph Swan.

Some of the first rotating machines used for electricity generation were based on water wheels and dynamos. However, water was not always available where power was needed and the trend among municipal power stations, the first important type of public power plant, was often to utilize steam engines and generators. These stations were initially built to provide electricity for lighting in cites. Early plants were generally small with a limited number of customers, but the area supplied by each power station gradually grew in size. At the same time there was little standardization and supply voltages varied from place to place and company to company. Meanwhile, there was an extended debate about the comparative merits of direct current and alternating current as the means of supplying electrical power. This was not resolved until well into the 20th century.

Lighting offered the first commercial use for electricity, but it proved an insufficient foundation for an industry. What accelerated the growth of electricity generation was its use for traction power, such as electric trams for urban transport and the underground railway systems in London and Paris. These were the kinds of projects that stimulated the construction of large power stations at the end of the 19th century and the start of the 20th century.

From here the industry spread rapidly, particularly with the use of electric motors in commerce and industry. The piecemeal development of the supply industry eventually became a problem and nationalization and standardization became common during the first half of the 20th century. Ironically, the first of these, nationalization, would be reversed in many countries during the last part of the same century. By that time electricity had become indispensable.

Although its origins are in the 19th century, few would dispute the argument that the growth of the electricity industry was a 20th-century phenomenon. There is little doubt, too, that by the end of the 21st century it will have become the world's most important source of energy. It is already starting to move into transportation with electric vehicles so that most types of energy needed can now be supplied electrically. It is worth remembering, however, that most of the key elements necessary for electricity generation, transmission, and distribution were developed during the 19th century.

EVOLUTION OF ELECTRICITY-GENERATING TECHNOLOGIES

The development of the electric power industry can be dated from the development of the dynamo or alternator. This allowed rotating machinery to be used to generate electricity. There were two sorts of generator used in the industry initially: the dynamo, which produced direct current, and the alternator, which produced and alternating current (the word "generator" can be used for both but it has become associated with the latter). The first practical dynamo was developed independently by Werner Siemens and Charles Wheatstone in 1867 and it was through the dynamo that the electric motor was discovered. However, the dynamo became displaced in most uses by the alternator, because alternating current distribution of power proved more efficient based on the technologies available at the time.

The first recorded power station appears to have been built in the Bavarian town of Ettal in 1878. This station used a steam engine to drive 24 dynamos, with the electricity used to provide lighting for a grotto in the gardens of the Linderhof Palace. Meanwhile, the first public power station was built in 1881 in Godalming in Surrey, United Kingdom. This station used two waterwheels to drive an alternator and provided power to two circuits—one at 250 V supplying power to 7 arc lights, and the second at 40 V providing power for 35 incandescent lamps.

As this brief historical snapshot demonstrates, both hydropower and steam power were already being used in the early days of the industry. Steam power was at this stage based on reciprocating steam engines, similar in concept to a piston engine. These engines were not ideal for the purpose because they could not easily develop the high rotational speeds needed to drive a generator effectively. This difficulty was eventually overcome with the invention of the steam turbine by Sir Charles Parsons in 1884. Fuel for these steam plants was usually coal, used to raise steam in a boiler.

Hydropower was an established source of mechanical power long before the steam engine was invented, so it was natural that it should provide one of the first engines used to drive dynamos and alternators. Water wheels were not the most efficient way of harnessing the power in flowing water but new turbine designs soon evolved. Much of the work on the main turbine types that are used today to capture power from flowing water—designs such as the Pelton and Francis turbines—was carried out in the second half of the 19th century.

By the beginning of the 20th century both the spark-ignition engine and the diesel engine had been developed. These too could be used to make electricity. Before World War II, work also began on the use of wind power as a way of generating power. Even so, steam turbine power stations burning coal, and sometimes oil or gas, together with hydropower stations provided the bulk of the global power generation capacity until the beginning of the 1960s.

In the 1950s the age of nuclear power was born. Once the principles were established, construction of nuclear power stations accelerated. Here, it was widely believed, was a modern source of energy for the modern age—it was cheap, clean, and technically exciting. Nuclear power continued to expand rapidly in the United States up to the late 1970s. In other parts of the world, uptake was less rapid but in western Europe, Great Britain, France, and Germany invested heavily, and in Scandinavia, Sweden developed a significant fleet of plants. In the Far East, Japan, Taiwan, and South Korea worked more slowly. Russia developed its own plants, which were used widely in eastern Europe, and India and China each began a nuclear program.

From the end of the 1970s the once-lustrous nuclear industry began to tarnish. Since then its progress has slowed dramatically, particularly in the west.

In Asia, however, the dream remains alive, although Japan's nuclear industry has been seriously damaged by the Fukushima disaster in 2011, the repercussions of which have reverberated around the world.

At the beginning of the same decade that saw nuclear fortunes turn, in 1973 to be precise, the Arab–Israeli war caused a major upheaval in world oil prices. These rose dramatically. By then oil had also become a major fuel for power stations. Countries that were burning it extensively began to seek new ways of generating electricity, and interest in renewable energy sources began to take off while the use of oil for power generation began to wane in all but the oil-producing countries of the Middle East.

The stimulus of rising oil prices led to the investigation of a wide variety of different alternative energy technologies, such as wave power, hot-rock geothermal power, and the use of ethanol derived from crops instead of petrol or oil. However, the main winners were solar power and wind power. Development took a long time, but by the end of the century solar and wind technologies had reached the stage where they were both technically and economically viable. By the end of the first decade of the 21st century both were growing strongly in overall installed capacity and, with prices coming down, this trend appears set to continue well into the century.

One further legacy of the early 1970s that began to be felt in the electricity industry during the 1980s was a widespread concern for the environment. This forced the industry to implement wide-ranging measures to reduce environmental emissions from fossil fuel–fired power plants. Other power generation technologies, such as hydropower, were affected too as their impact on local environments and people were reassessed.

The gas turbine began to make a major impact during the 1980s as an engine for power stations. The machine was perfected during and after World War II as an aviation power unit but soon transferred to the power industry for use in power plants supplying peak demand. During the 1980s the first large base-load power stations using both gas turbines and steam turbines in a configuration known as the combined cycle plant were built. This configuration has become the main source of new base-load capacity in many countries where natural gas is readily available.

The first years of the 21st century have seen increased emphasis on new and renewable sources of electricity. Fuel cells, a technically advanced but expensive source of electricity, are approaching commercial viability. There is renewed interest in deriving energy from ocean waves and currents, and from the heat in tropical seas. Offshore wind farms have started to multiply around the shores of Europe.

The story of power generation across the 21st century is likely to be the contest between these new technologies and the old combustion technologies for dominance within the power generation industry. And while they battle for supremacy there remains one technology—nuclear fusion—that has yet to prove itself but just might sweep the board.

EVOLUTION OF ELECTRICITY NETWORKS

For electricity from a power station or power-generating unit to be delivered to a customer, the two must be connected by an electricity network. Over the past century these networks have developed into massive systems.

When the industry was in its infancy, networks were a simple pattern of lines radiating from a power station to the small number of customers that each power station supplied, usually with a number of customers on each line. When the number of customers was small and the distances over which electricity was transported were short, these lines could operate with either direct current (DC) or alternating current (AC).

As the distances increased, it became necessary to raise the voltage at which the electricity was transmitted to reduce the current and the resistive losses in the lines when high currents were flowing. The AC transformer allowed the voltage on an AC line to be increased and then decreased again efficiently and with relatively ease, whereas this was not possible for the DC system. As a consequence, alternating current became the standard for most electricity networks.

Alternating current continued to dominate across the 20th century, but developments in power electronics led to a resurgence in interest in the DC transmission of power at the end of the century in the form of high-voltage DC lines. These are increasingly used for sending large amounts of power over long distances for which they are proving more efficient than conventional AC lines.

Back at the start of the 20th century, the growth in size of what was initially a myriad of independent electricity networks soon led to overlap between service areas. While competition was good for the electricity market, the range of different operating standards, particularly voltages and frequencies, made actual competition difficult. A proliferation of independent networks was also costly, and in the final analysis it was unnecessary because if different operators standardized on their voltages, the suppliers of electric power could all use the same network rather than each building its own.

Standardization was pushed through in many countries during the first half of the 20th century and national grid systems were established that were either government owned or controlled by legislation to ensure that the monopoly they created could not be exploited. However, there are still vestiges of the early market proliferation of standards to be found today in regional variations, such as the delivery of alternating current at either 50 Hz or 60 Hz and the different standard voltage levels used.

As national networks were built up, a hierarchical structure became established based on the industry model in which electric power was generated in large central power stations. These large power plants fed their power into what is now the transmission network, a high-voltage backbone that carries electricity at high voltage from region to region. From this transmission network, power is fed into lower-voltage distribution networks and these then deliver the power to the customers.

An electricity network of any type must be kept in balance if voltage and frequency conditions are to be maintained at a stable level. This is a consequence of the ephemeral nature of electricity. The balance between the actual demand for electricity on the network and the power being fed into it must be maintained within narrow limits. Any deviation from balance leads to changes in frequency and voltage and, if these become too large, can lead to a system failure.

The organization charged with maintaining the balance is called the system operator. This organization has limited control over the demand level but it must be able to control the output of the power plants connected to its network. For most networks this has traditionally involved having a variety of different types of power stations supplying power. The first of these are base-load power plants. These are usually large fossil fuel and nuclear power plants (but they may also include hydropower) that keep running at maximum output all the time, supplying the basic demand on the network. Next are intermediate-load power plants, often gas turbine based, which do not run all the time but might start up in the morning to meet the daytime rise in demand and then close down in the evening when demand begins to fall again. These two types can supply the broad level of demand during both day and night but there will always be a need for even faster-acting plants that can provide the power to meet sudden peaks in demand. These are called peak-load or peaking power plants. In general, the power from base-load power plants is the cheapest available, that from intermediate load plants is more expensive, and that from peak-load plants is the most expensive.

RENEWABLE ENERGY AND DISTRIBUTED GENERATION

Most forms of renewable energy do not fit happily into this operational structure. When electricity from them is available it must be used as if it were from a base-load plant, but because their output is intermittent, they cannot be relied on in the same way as a base-load plant. As a consequence, the introduction of large quantities of renewable energy into electricity systems that began at the beginning of the 21st century is leading to important changes in the way grid stability is maintained.

While the output of a conventional power plant will, barring accidents, remain steady and predictable, many renewable sources including wind and solar energy are, as just noted, intermittent and often unpredictable. This means that system operators must now manage not only variable demand from customers but also a variable supply. New strategies are being adopted by the system operators, including the use of highly detailed weather forecasting, so that the output from variable renewable plants can be predicted ahead of time and alternative capacity arranged where a shortfall is expected. The amount of electricity storage capacity included on a grid is expected to grow too as a means of balancing renewable generation.

The growth in the use of solar and wind power has led to other changes. One of these is an increase in the number of power stations that feed power into the distribution level of the grid rather than the transmission system level. This has arisen partly because many renewable power plants are too small to provide power into the backbone of the grid network. However, another important factor is the increased use of electricity generated locally.

Local generation might be a rooftop solar panel supplying power to a single domestic household or a small wind farm providing electricity to the community that owns it. In both cases, the electric power never enters the transmission backbone but travels from generator to consumer either internally or across the distribution system. Distributed generation, as generation at this point in the grid hierarchy is called, is valuable because it means that power does not have to be transported far between the generator and the user, thereby reducing transmission losses. In the extreme case, such as power from a rooftop solar panel being used in a domestic household, virtually all transportation losses are eliminated.

The use of distributed generation allows new capacity to be added in small tranches as demand grows in a particular area. However, it creates difficulties for distribution networks. Traditionally these have been designed as passive networks that deliver power from transmission networks to customers. They do not have the ability to balance their own network power locally because it has never been necessary for them to have this capability. Distributed generation alters the situation, and it is leading to the development of balancing capability at the distribution level. It is likely that future electricity networks will resemble an interconnection of such semi-autonomous distribution networks rather than a single hierarchical network.

One of the technologies that will enable this to take place is called the smart grid. A smart grid is an electricity network in which the system operator, the power suppliers, and the power consumers can all interact and communicate in real time. To achieve this it is necessary to build a computer network that runs in parallel to the electricity network. This type of system is beginning to appear in parts of the world and it is likely to become common over the next 10–20 years.

The two-way communication across the network will eventually allow network system operators to control electricity-using devices such as washing machines or water heaters in homes, as well as consumption in commercial and industrial organizations, in response to changes in demand. This type of flexibility will help integrate larger amounts of renewable energy into grids without endangering stability, as well as making electricity use more efficient.

A BRIEF POLITICAL DIVERSION

During the last years of the 19th century, when the technology was in its infancy, the generation of electricity was seen as one more opportunity for entrepreneurs and joint stock companies to make money. After all, electricity was not unique. There were other means of delivering energy; district heating was already common in the United States and in some European cities, while hydraulic power was sold commercially in cities like London. Gas, produced from coal, was also sold for lighting and other commercial uses. As a result, the early history of the electricity industry was one of small, privately owned companies competing with other energy companies. Gradually, however, the distribution of electricity rendered most other ways of distributing energy across a network obsolete.

In the 20th century, as the primacy of electricity became obvious, the distribution of electricity gradually came to be seen as a public service—that is, a public utility. Like water, sewage, and, later, the gas supply, electricity was needed to operate a modern civilization. In much of the world, the electricity industry was absorbed by government and became publicly owned. In countries such as the United States where this did not happen, legislation was introduced to govern the supply. In the vast majority of countries, however, national, vertically integrated utilities were established. Even in the United States, large utilities generated, transmitted, and distributed power.

Nationalization was a pragmatic solution, but it could also be seen as part of a socialist or left-leaning ideology that favored central government over independent business. In the late 20th century, western political ideologies began to shift, particularly among the more right-leaning political parties. Government ownership of industry, including the electricity industry, began to be seen as unnecessary and uneconomic. A move began to convert publicly owned utilities into privately held companies. Alongside this, utility legislation was relaxed to open electricity markets to competition.

Returning a national electricity company to private ownership was not easy. Power stations had to be sold individually, to a number of different companies, to create market competition. Transmission and distribution were natural monopolies; these too were sold to private companies, but these companies had to operate within rules set by the government. Then a complex electricity market had to be established in which the competing generating companies could sell their electricity.

In spite of the complexity, by the beginning of the 21st century liberalization had become a global phenomenon. A few centralized governments still retained full control over their electricity industries, but most paid at least lip-service to the concept of liberalization. The result was both successes and failures. California, for example, recorded a dramatic failure when liberalization resulted in a virtual breakdown of its electricity supply system, with almost catastrophic consequences. The cost of electricity in California rose dramatically as a result. Elsewhere prices fell after liberalization.

Private sector ownership of the electricity industry is now the predominant global model. However, it is not without its drawbacks. If state control of the electricity industry was seen to be overbearing and too rigid, a liberalized industry can seem to have too much freedom. Economic rather than political considerations become paramount.

In the United Kingdom during the 1990s the result of liberalization of the electricity system was a rush to build gas-fired power stations because they were cheap and, when natural gas was cheap, they were the most economical source of new generating capacity. However, when gas prices rose dramatically, these plants became uneconomical, and by the end of the decade plants were being taken out of service. A centrally planned electricity supply system would probably have adopted a more balanced and diverse approach to increasing generating capacity.

Liberalization also makes government energy policy more difficult to implement. Renewable energy offers a good example. A government that wants to increase the proportion of electricity generated from renewable sources cannot simply pass an order down the line to its national utility. It must use taxes and systems of allowances and penalties to encourage energy suppliers to make the choices it wants. However, generating companies may simply chose to pay the penalties if that is the most economically attractive option. In that case the desire of government is ignored.

Given such difficulties, it is arguable that the balance has swung too far in favor of market forces in the electricity sector. Competition, the whole basis for market economics, can be extremely limited in some of these electricity markets. The structures that need to be created to turn what looks from most angles to be a natural monopoly—electricity transmission and distribution, for example—into a competitive industry can be artificial and convoluted. In spite of such criticisms, the model has become entrenched.

It seems likely that free-market rules will continue to dominate the electricity industry, but it would be foolish to predict that it will be so. It is far from uncommon for one generation to reverse the policies of a previous generation. However, there is no sign of any movement in the opposite direction in the second decade of the 21st century.

SIZE OF THE INDUSTRY

How big is the global electricity industry? Tables 1.1 and 1.2 provide the answer, at least as far as power generation is concerned. Table 1.1 shows the global generating capacity in 2010 broken down by power station type together with projections for the sizes in 2020, 2030, and 2040.

Total global generating capacity in 2010 was 5061 GW. Based on the predictions contained in Table 1.1, this capacity will rise to 6221 GW in 2020, 7214 GW in 2030, and 8254 GW in 2040. This is more than twice the capacity in

TABLE 1.1 Global Generating Capacity by Source (GW)

	2010	2020	2030	2040
Liquid fuels	399	354	311	283
Natural gas	1299	1434	1699	2057
Coal	1649	1884	2127	2300
Nuclear	381	490	630	717
Hydropower	917	1190	1363	1619
Wind	183	480	619	726
Geothermal	10	18	23	29
Solar	37	157	204	266
Other renewable sources	187	215	238	256
Total	5061	6221	7214	8254

Source: *International Energy Outlook 2013,* U.S. Energy Information Administration.

TABLE 1.2 Global Net Electricity Generation by Source (TWh)

	2010	2020	2030	2040
Liquid fuels	914	822	746	678
Natural gas	4479	5541	7206	9372
Coal	8052	10,122	12,304	13,891
Nuclear	2620	3638	4755	5492
Hydropower	3402	4452	5177	6232
Wind	342	1136	1544	1839
Geothermal	66	133	171	220
Solar	34	240	327	452
Other renewable sources	332	549	729	858
Total	20,240	26,632	32,959	39,034

Source: *International Energy Outlook 2013,* U.S. Energy Information Administration.

2000, which stood at 3366 GW. These predictions depend on a number of assumptions about the way in which the industry will develop and, in particular, how active governments are promoting renewable energy. Growth could plausibly be either slower or faster but the figures in Table 1.1 offer a median projection.

When the figures are broken down by the type of generating station or energy source, coal can be seen to be the largest provider of power in terms of generating capacity, accounting for just under 33% of the global total. This continues to be the case in the predictions for future years shown. Natural gas is the second most important energy provider, followed by hydropower. Between them, these three accounted for 76% of global capacity in 2010. This is predicted to fall to 72% by 2040.

Aside from hydropower, wind and solar power are considered the most important renewable sources today, and these are the generating technologies that are showing the most rapid growth in global capacity. Both are expected to show strong growth during the period covered by Table 1.1, and it is possible that the actual growth, particularly for solar power, could outstrip these predictions.

A regional breakdown of capacity and capacity growth from the same source as these figures shows that the fastest growth in generating capacity is taking place among developing countries where average annual growth is expected to be 2.3% between 2010 and 2040. Within developed countries, on the other hand, the predicted rate of growth, at 0.9%, is notably slower.

Table 1.2 shows parallel figures to those shown in Table 1.1 but for global net electricity generation, broken down by type. The total amount of electricity generated in 2010 was 20,240 TWh. As before the table contains predictions for the production level in 2020, 2030, and 2040. These figures show that production is expected to virtually double between 2010 and 2040, when it is predicted to reach 39,034 GWh. This is 167% more than the net generation in 2000, when global output stood at 14,618 TWh.

As might be expected from Table 1.1, the most important global source of electricity is coal, which provided 8052 TWh in 2010, or just under 40% of the total. This is expected to decline to 35% by 2040, but it will remain the single most important source of electricity based on these predictions. Natural gas is the second largest source in 2010, providing 4479 TWh or 22% of the total. Hydropower is the third most significant source with 17% of the total. Once again these three supply the bulk of the global electricity supply—79% of the total in 2010.

It is notable, when comparing Tables 1.1 and 1.2, that while nuclear generating capacity accounted for only 8% of total generating capacity in 2010, it actually supplied 13% of total output, proportionally much larger than the contribution from hydropower. This is a reflection of the fact that while nuclear power plants generally operate as base-load power stations, generating at close to full power for most of the time, hydropower plants provide a variable output depending on the season. It is also becoming increasingly common for hydropower plants to be used for grid support, providing power when other sources cannot. This means that they often provide less power than they are technically capable of supplying.

The output from both solar and wind power plants is also lower than that from fossil fuel and nuclear power plants of similar generating capacity. This is because the energy sources—that is, the wind and sun—are intermittent and do not provide energy all the time. While a coal-fired power plant or a nuclear power plant may be able to generate power for 80–90% of the time, a wind plant will be performing well if it can generate the equivalent of full power for 40% of the time.

While the figures in the Tables 1.1 and 1.2 provide a broad picture of production and capacity globally, there are important regional differences that they cannot convey. These differences are a reflection of the resources available in different parts of the world. For example, many South American countries rely heavily on hydropower, whereas the main industrial nations in Asia, China, and India are heavily dependent on coal power. The United States also provides a large part of its electricity from coal. In contrast, coal use in Europe is much lower, and this region as a whole has the largest new renewable generating capacity, mainly wind and solar power. (Hydropower is considered an old renewable technology.)

The electricity industry is large and it is also conservative. Change is generally slow and takes place over generations. Most power stations have a lifetime of 30 years, some much longer, and capacity is only replaced when it is exhausted. Nevertheless, the composition of the world's generating capacity will not remain static and new technologies will gradually displace older ones. If the 20th century was the century when coal and hydropower dominated, the 21st century can be expected to see much more renewable generation. Even so, as Table 1.2 shows, coal and hydropower will remain important well into the century.

Chapter 2

Electricity Generation and the Environment

The power generation industry, taken as a whole, is the world's biggest industry and it has the largest impact of any industry on the environmental conditions on Earth. Some of the effects caused by power generation, particularly the ones associated with the combustion of fossil fuels, are far-reaching both geographically and temporally. However, all types of power generation will have some effect.

In the early days of the industry, environmental considerations were rarely taken into account when power stations were built. Economic considerations were the priority and an awareness of the dangers associated with power generation was slow to register. It was the effect of air pollution created by coal combustion on human health that provided one of the first warning signs. Early mitigating measures to combat this included the use of tall stacks to release power plant flue gases higher into the atmosphere.

A much greater awareness of the global environment, and of humankind's impact on it, began to bloom during the 1970s. From the power industry's point of view this was a turning point, too. Perhaps it was simply a matter of the size of the global power sector by that time, but since then the history of the industry has been punctuated by graphic illustrations of the problems it can produce. The damaging effects of acid rain were recognized during the 1980s. A nuclear disaster at Chernobyl in the Ukraine in 1986 illustrated the dangers of nuclear power. During the 1980s and 1990s there were critical reviews of a range of large hydropower plants that eventually prompted a change in the way such schemes were evaluated. Then, during the 1990s, the dangers of global warming were recognized and the burning of fossil fuels was identified as one of the probable causes.

In the 21st century, environmental considerations are shaping the way in which the industry is evolving. The emissions from fossil fuel power plants are being radically reduced, and although this does not yet include carbon dioxide, it is being targeted for future removal. Alongside this, renewable energy is being promoted as a cleaner, more sustainable means of generating electricity. This, in turn, is changing the way electricity networks operate.

These changes will continue throughout the century so that by the end of the century the electricity-supply industry may well be unrecognizable to somebody used to the configuration today. But these changes, too, will have environmental consequences. Already there are signs of what is to come with many more wind farms and large area solar power plants found both in the landscape and offshore. Such changes will inevitably lead to new problems. An industry as significant as power generation cannot avoid being in conflict, at one level or another, with the global environment.

EVOLUTION OF ENVIRONMENTAL AWARENESS

Humans have always changed their surroundings. Some of those changes we no longer even recognize—for example, the clearing of forests to create Europe's agricultural farmlands. No one now sees these fields as forests that once were. Similar changes elsewhere are more obviously detrimental to local or global conditions. Tropical rain forests grow in the poorest of soils. Clear them and the ground is of very little use. Not only that but the removal of forest cover can lead to erosion and flooding, as well as the loss of ground water. Most of these effects are negative.

Part of the problem is the ever-increasing size of the human population. Where native tribes could survive in the rain forests in Brazil, the encroachment of outsiders has led to their erosion. A similar effect is at work in power generation. When the demand for electricity was limited, the effect of the few power stations needed to supply that demand was small. But as demand has risen, so has the cumulative effect. Today that effect is of such a magnitude that it cannot be ignored.

Consumption of fossil fuels is the prime example. Consumption of coal has grown steadily since the Industrial Revolution. The first sign of trouble resulting from this practice was the ever-worsening pollution in some major cities. In London the word "smog" was invented at the beginning of the 20th century to describe the terrible clouds of fog and smoke that could remain for days. Yet it was only in the 1950s that legislation was finally introduced to control the burning of coal in the U.K. capital.

Consumption of coal still increased, but with the use of smokeless fuel in cities and tall stacks outside, problems associated with its combustion appeared to have been solved. Until, that is, it was discovered that forests in parts of northern Europe and North America were dying and lakes were becoming lifeless. During the 1980s the cause was identified: acid rain resulting from coal combustion. More legislation, aimed at controlling the emissions of acidic gases such as sulfur dioxide and nitrogen oxide, was introduced.

Acid rain was dangerous but worse was to come. By the end of the 1980s scientists began to fear that the temperature on the surface of the Earth was gradually rising. This has the potential to change conditions everywhere. Was this a natural change or human-made? Scientists did not know.

As studies continued, evidence suggested that the effect was, in part at least, human-made. The rise in temperature followed a rise in the concentration of some gases in the atmosphere. Chief among these was carbon dioxide. One of the main sources of extra carbon dioxide was the combustion of fossil fuels such as coal.

If this is indeed the culprit, and the weight of evidence available makes it prudent to assume that it is, then consumption of fossil fuels must fall or measures must be introduced to remove and secure the carbon dioxide produced. Otherwise, the global temperature is likely to rise to a level that could cause disruption and destruction in many parts of the world. In the worst case it is possible to imagine some catastrophic change to global conditions. It has now become one of the main challenges for governments all over the world to reduce the amount of carbon dioxide being released into the atmosphere without crippling their economies.

The way in which fossil fuels are used in power generation is gradually changing as a result of these discoveries and the legislation that has accompanied them. Other technologies also face challenges. Nuclear power is considered by some to be as threatening as fossil fuel combustion, though it has its advocates too. Hydropower has attracted bad publicity in recent years but should still have an important part to play in future power generation. Meanwhile, there are individuals and groups prepared to go to almost any lengths to prevent the construction of wind farms, which they consider unsightly, and objections to solar power plants have started to be heard.

At the same time, electricity is vital to modern living. Therefore, unless the world is going to regress, technically, the supply of electricity must continue and grow. On that basis, compromises must be sought and technical solutions found that do not result in irrevocable damage. These are the challenges that the power industry faces, and with it the world.

ENVIRONMENTAL EFFECTS OF POWER GENERATION

Much human activity has an effect on the environment and, as already outlined, power generation is no exception. Some of these effects are more serious than others. The atmospheric pollution resulting from coal, oil, and gas combustion has had obvious effects. But combustion of fossil fuel also releases a significant amount of heat into the environment, mostly as a result of the inefficiency of the energy conversion process. Is this a serious side effect? In most cases, it probably is not.

Power stations have a physical presence in the environment. Some people will consider this a visual intrusion. Most make noises, another source of irritation. There are electromagnetic fields associated with the passage of alternating currents through power cables. A power plant needs maintaining, servicing, and often to be provided with fuel, which will generate road or rail traffic. All of these factors will have an impact on people living within the vicinity of a facility even if they do not affect a wider area.

Clearly, some of these effects are more far-reaching than others. Even if they are not far-reaching, the local effects of a power station may be a significant issue for a population that lives immediately adjacent to the facility. Deciding what weight must be given to such considerations when planning future generating capacity can be a fearsomely difficult issue. It is the big issues, however, particularly global warming, that will have the most significant effect on the future of power generation.

CARBON CYCLE AND ATMOSPHERIC WARMING

The combustion of fossil fuels such as coal, oil, and natural gas releases significant quantities of carbon dioxide into the atmosphere. These fuels were created from organic material growing on Earth's surface that became buried, locking the carbon they contained within the body of the planet. Since the Industrial Revolution the use of these fuels has accelerated. The consequence appears to have been a gradual but accelerating increase in the concentration of carbon dioxide within Earth's atmosphere.

Before the Industrial Revolution the concentration of carbon dioxide in the Earth's atmosphere was around 270–280 parts per million (ppm). Between 1700 and 1900 there was a gradual increase in atmospheric concentrations but from 1900 onwards, the concentration changed more rapidly, as shown in Table 2.1. From 1900 to 1940 atmospheric carbon dioxide increased by around 10 ppm,

TABLE 2.1 Atmospheric Carbon Dioxide Concentrations*

Decade	Carbon Dioxide Concentration (ppm)
1700	270–280
1900	293
1940	307
1960	312
1970	326
1980	339
1990	354
2000	369
2010	390
2050	440–500
2100	500–700

**Data before 1959 is derived from ice core measurements. Data since 1959 is based on measurements at Manua Loa in Hawaii by Dr. Pieter Tans, NOAA/ESRL, and Dr. Ralph Keeling, Scripps Institution of Oceanography. Predictions are based on generally proposed levels from different sources.*

Source: U.S. Earth System Research Laboratory.

from 1940 to 1980 it increased by 32 ppm, and by 2000 it had increased by a further 30 ppm. In 2010 the concentration was 21 ppm higher than in 2000. By then the total concentration was 390 ppm, around 40% higher than in 1700.

If the increase in carbon dioxide concentration is a direct result of the combustion of fossil fuels, then it will continue to rise until that combustion is curbed. Estimates of future concentrations are at best speculative, but Table 2.1 includes a range of estimates for both 2050 and 2100. The worst case in the table shows concentrations doubling in 100 years.

While the increase in carbon dioxide concentration is clear and continuous, the increase in global temperature is more variable. The evidence for a fossil fuel connection with the increase in concentration of carbon dioxide is compelling, but the cycling of carbon among the atmosphere, sea, and biosphere is so complex that it is impossible to be certain how significant the human-made changes are, or what other factors may be involved.

The atmospheric emissions of carbon from human activities such as the combustion of coal, oil, and natural gas amounted to a total of around 8.6 Gtonnes in 2012. While this is an enormous figure, it is tiny compared to the total carbon content of 750 Gtonnes in the atmosphere. This atmospheric carbon is part of the global carbon cycle. There are roughly 2200 Gtonnes of carbon contained in vegetation, soil, and other organic material on Earth's surface, 1000 Gtonnes in the ocean surfaces, and 38,000 Gtonnes in the deep oceans.

The carbon in the atmosphere, primarily in the form of carbon dioxide, is not static. Plants absorb atmospheric carbon dioxide during photosynthesis, using the carbon as a building block for new molecules. Plant and animal respiration, on the other hand, part of a natural process of converting fuel into energy, releases carbon dioxide to the atmosphere. As a result there are probably around 60 Gtonnes of carbon cycled between vegetation and the atmosphere each year, while an additional 100 Gtonnes are cycled between the oceans and the atmosphere by a process of release and reabsorption. Thus, the cycling of carbon between the atmosphere and Earth's surface is a complex exchange into which the human contribution from fossil fuel combustion is small.

The actual significance of the additional release of carbon dioxide resulting from human activity depends on the interpretation of various scientific observations. The most serious of these relate to a slow increase in temperature at Earth's surface. This has been attributed to the greenhouse effect, whereby carbon dioxide and other gases in the atmosphere allow the sun's radiation to penetrate the atmosphere but prevent heat leaving, in effect acting as a global insulator.

If human activity is responsible for global warming, then unless carbon dioxide emissions are controlled and eventually reduced, the temperature rise will continue and probably accelerate. This will lead to a number of major changes to conditions around the globe. The polar ice caps and glaciers will melt, leading to rises in sea level, which will inundate many low-lying areas of land. Climate conditions will change. Plants will grow more quickly in a carbon dioxide–rich atmosphere.

Not all scientists agree that changes in our practices can control the global changes. There have been large changes in atmospheric carbon dioxide concentrations in the past, and large temperature swings. It remains plausible, though unlikely given the weight of evidence now available, that both carbon dioxide concentration changes and global temperature changes are part of a natural cycle and that the human contribution has little influence.

In the second decade of the 21st century the weight of the scientific evidence suggests a strong link between human release of carbon dioxide and global warming, but it may be impossible to find absolutely conclusive proof. Unfortunately, the nature of scientific inquiry will always leave open the possibility, however remote, of an alternative interpretation. But in the meantime conditions will continue to change. And if human activity is responsible, the change may eventually become irreversible.

Even if a link between an increase in carbon dioxide concentration and the rising global temperature cannot be made with absolute certainty, it is clear that combustion of fossil fuel is creating more carbon dioxide than would naturally have entered the atmosphere. This fact itself provides a powerful reason to act to slow and eventually reverse the release of the gas. However, there are also economic considerations and these are generally not in favor of any rapid adjustment. How the competing demands are balanced is likely to determine the fate of the globe.

CONTROLLING CARBON DIOXIDE

Fossil fuels are all derived from trees and vegetation that grew millions of years ago and subsequently became buried beneath the surface of the earth. Without man's intervention the carbon contained in these materials would have remained buried and removed from the carbon cycle. As a result of human activity they have been returned to the carbon cycle.

An immediate cessation of all combustion of fossil fuel would stabilize the situation. That is currently impossible. Too much global economic activity depends on burning coal and gas. As the figures in Table 1.2 showed, predictions suggest that the use of fossil fuels, particularly coal, will increase over the next decades, not decrease. The popular strategy in some regions of switching fuel from coal to gas reduces the amount of carbon dioxide generated but does not eliminate it.

One short-term measure would be to capture the carbon dioxide produced by a combustion power station and store, or sequester, it in a way that would prevent it from ever entering the atmosphere. Technologies exist that are capable of capturing the carbon dioxide from the flue gas of a power plant, and these are being developed for commercial deployment. Finding somewhere to store it poses a more difficult problem.

One solution is to pump it into exhausted oil and gas fields. There are other underground strata in which it might be stored. A third possibility is to store it at the bottom of the world's oceans. The enormous pressures found there would solidify the gas and the solid would remain isolated unless disturbed. Whether this would be environmentally acceptable is another matter.

These solutions are all expensive and none is particularly attractive. However, they may become necessary as short-term solutions. Over the longer term the replacement of fossil fuels with either renewable technologies that do not rely on combustion or with biomass-generated fuel that releases carbon dioxide when burned but absorbs it again when it is regrown, will be necessary if conditions are to be stabilized. That appears likely to take most of the coming century, at least.

HYDROGEN ECONOMY

A switch to sustainable renewable technologies would appear to offer a practical means to control power plant emissions of carbon dioxide, but it will not solve all global problems associated with fossil fuel. What about all the other uses, particularly for automotive power? A more radical solution might be to switch from an economy based on fossil fuel to one based on hydrogen.

Fossil fuels, particularly oil and gas, have become a lynch pin of the global economy because they are so versatile. The fuels are easily stored and transported from one location to another. They can be used in many different ways, too: power stations, internal combustion engines, cookers, refrigerators—all these and more can be powered will one of these fuels.

Renewable electricity sources such as hydropower, solar power, wind power, and biomass can replace fossil fuel in power generation but they cannot easily be adapted to meet all the other uses to which fossil fuel is put. The most salient of these is transportation. One solution that is being followed by vehicle manufacturers is to build electric vehicles that have batteries to provide their energy source. These batteries must then be recharged regularly from the grid, perhaps using renewably generated electricity. This is one vision of a fossil fuel–free future.

Hydrogen offers an alternative. Instead of a battery, a vehicle can carry a supply of hydrogen that it can burn in a conventional reciprocating engine to provide power. Alternatively, the hydrogen can be used to provide energy for a fuel cell that will generate electricity from it. This then provides convergence with battery-powered vehicles.

Hydrogen has the immense advantage that it can replace fossil fuels not just in vehicles but in virtually all applications. Not only can it be used to power an internal combustion engine, but it can be burned to provide heating or cooling. Conventional fossil fuel power plants can burn it to generate electricity. Moreover, it can be stored and transported with relative ease. And it is clean. When it is burned, the only product of its combustion is water.

Where would the hydrogen for a hydrogen economy be found? The primary source would be water, and the best way of making it would be by use of electrolysis. Renewable energy power plants would generate electricity and the electricity would be used to turn water into hydrogen and oxygen. The hydrogen would be captured and stored for future use.

This may seem like an expensive and inefficient method of generating fuel. It is, although scientists are working hard to improve the efficiency. For a hydrogen economy to work today electricity from renewable sources needs to become much cheaper—cheaper probably than all but the cheapest electricity today. Even so, it offers a vision for the future in which life continues in much the same way as it does today. That can be seductive.

ECONOMICS OF ELECTRICITY PRODUCTION

What, exactly, is the cost of a unit of electricity today? That is not an easy question to answer. In basic economic terms the cost depends on the cost of the power station—how much it costs to build (this figure should include the cost of any loans needed to finance the construction and this can end up being one of the main contributors to the final cost of the electricity)—the cost of operating and maintaining it over its lifetime, which is typically 30 years for a combustion plant; and the cost of fuel. If all these numbers are added up and divided by the number of units of electricity the plant produces over its lifetime, then this is the basic unit cost of electricity. In most cases there will then be an addition to cover either profits if the plant is owned by the private sector or for future investment if the plant is publicly owned, to arrive at the cost to the consumer of the power.

With some adjustment for the change in the value of money over the lifetime of the plant, this calculation results in a figure called the levelized cost of electricity (LCOE) for a particular plant or technology, a figure that is often used to compare the value of different types of power plants.

The LCOE calculation is theoretical. It is impossible to know in advance exactly how a power plant will perform economically, because future conditions and costs are impossible to obtain in advance. On the other hand, anybody building a facility will want to know if it is going to be economical before actually constructing it. To get around this problem, guesses have to be made about the future performance and costs over its lifetime. Assumptions must also be made about some economic factors that take account of the changing value of money. These add a great deal of uncertainty to the result but the basic equation remains the same.

When this calculation is carried out for a range of different types of generating plants, the cheapest new source of electricity in many developed regions of the world will prove to be a natural gas–fired combined-cycle power plant. This type of plant is cheap, quick to build, and relatively easy to maintain. Capital cost of construction is low but the fuel price is relatively high, so for a plant of this type the cost of fuel is the most significant determinant of electricity prices. While gas is cheap, so is electricity. However, when gas prices rise, the plant can rapidly become uneconomical compared to other sources.

Gas prices have a history of volatility and this adds an element of risk when calculating the cost of electricity from a plant of this type. Other power stations, particularly those that use renewable energy where the energy source (the fuel) is free, are not exposed to this volatility. So while the cost of electricity from a plant of this type may be on average more expensive, it is also more stable. This is being recognized as a significant factor when evaluating the economic viability of new sources of electric power.

EXTERNALITIES

Risk and prices volatility aside, does the basic economic equation described before take account of all the factors involved in generating electricity? There is a growing body of opinion that says no. It says that there are other very important factors that need to be taken into account too. These are generally factors such as the effect of power production on the environment and on human health, factors that society pays for but not the electricity producer or consumer directly. These factors are called externalities.

A major study carried out by the European Union (EU) with support from the United States over a decade in the 1990s estimated that the cost of these externalities within the EU, excluding the cost of global warming, were equivalent to between 1% and 2% of EU gross domestic product.

The cost of electricity in the EU in 2001, when the report of the study was published, was around €40/MWh. External costs for a variety of traditional and renewable energy technologies as determined by the study are shown in Table 2.2. Actual external costs vary from country to country and the table

TABLE 2.2 External Cost of Power Generation Technologies

	External Cost (€/MWh)
Coal and lignite	20–150
Peat	20–30
Oil	20–111
Gas	10–40
Nuclear	2–7
Hydropower	0–10
Biomass	2–30
Solar photovoltaic	6
Wind	1–3

Source: Figures are from the ExternE project funded by the European Union and United States.

shows the best and worst figures across all countries. These figures indicate that coal combustion costs at least an additional €20/MWh on top of the €40/MWh paid by the consumer, and could cost as much as an additional €150/MWh. Gas-fired generation costs at least an additional €10/MWh, while the external costs for nuclear and most renewable technologies were a fraction of this.

If consumers were forced to pay these external costs—by the imposition by governments of some form of surcharge, for example—the balance in the basic equation to determine the cost of electricity would shift in favor of all the non-combustion technologies. Of course, such a move would initially penalize consumers because a high proportion of the world's electricity comes from fossil fuel and the capacity cannot be replaced overnight. It would not suit fossil fuel producers either and it would drive up manufacturing prices, initially at least, affecting global economics. Over the long term, however, if the analysis is correct, everyone would benefit.

LIFE-CYCLE ASSESSMENT

Another important tool for establishing the relative performance of power generation technologies is a life-cycle assessment. The aim of a life-cycle assessment is to measure the performance of a power plant with reference to one or more parameters, such as its emissions of carbon dioxide or the energy efficiency of its power generation. The assessment covers the complete life cycle of the plant starting from the manufacture of the components that were used to construct it and ending with its decommissioning.

The LCOE calculation discussed earlier is a type of life-cycle assessment of the economic performance of a power station. If one were instead examining the lifetime amount of carbon dioxide produced by a fossil fuel–fired power plant, one would examine not only the amount produced by burning the fuel in the plant, but also that produced when electricity or some other form of energy was used to manufacture the components used to build the plant, and any produced when the plant was dismantled and recycled. All these quantities can then be added up and divided by the amount of electricity the plant generates to provide a figure for the amount of carbon dioxide for each unit of power.

When carbon dioxide emissions are studied figures show, as would be expected, that coal-, gas-, and oil-fired power plants release massively more carbon dioxide for each unit of electricity they produce than do most renewable technologies. Typical figures are given in Table 2.3. Similar results are found for other common combustion plant emissions such as sulfur dioxide, carbon monoxide, and nitrogen oxide.

There are other types of life-cycle assessment that can be carried out. Another of relevance to power plant performance is the total energy balance of a plant. This involves a calculation of the number of units of energy a power station produces for each unit of energy it consumes over its lifetime. Energy is used to manufacture components for a power plant. For a fossil fuel plant,

TABLE 2.3 Lifetime Emissions of Carbon Dioxide for Various Power Generation Technologies*

	Carbon Dioxide Emissions (tonnes/GWh)
Coal	964
Oil	726
Gas	484
Nuclear	8
Wind	7
Photovoltaic	5
Large hydropower	4
Solar thermal	3
Sustainable wood	−160

**This is a European Commission–supported report published by the Delft University Wind Energy Research Institute.*

Source: *Concerted Action for Offshore Wind Energy in Europe*, European Union, 2001.

energy is used to harvest and deliver the combustion fuel to a power plant. The fuel itself contains energy that is consumed. And energy is consumed during the decommissioning of a power plant. All these units of energy must be added together and then divided by the total number of units of energy the power station delivers.

When the energy content of the fuel is included in the calculation, renewable generating technologies often appear as relatively poor performers. This is because while a thermal power plant can convert between 45% and 60% of the energy in its fuel into electricity, a wind turbine will generally struggle to convert more than 40% of the wind energy into electrical energy and a solar plant typically less than 20%. On the other hand, while these renewable efficiencies are relatively low, the energy that is not converted is not wasted—it simply isn't used. The energy from a combustion fuel that is not converted into electricity is released as waste heat.

A more useful comparison can be made by excluding the actual energy content of the energy source. One set of energy payback ratios calculated on this basis is shown in Table 2.4. The figures in the table confirm that when the energy source is excluded most renewable technologies score more highly than fossil fuel power plants on this measure, though the manufacture of solar cells is relatively energy intensive, making the energy payback ratio for this technology low. Some newer solar photovoltaic fabrication technologies will have better

TABLE 2.4 Energy Payback Ratio for Various Power Generation Technologies

	Energy Payback Ratio
Coal	3:5
Oil	1:3
Gas	3:5
Nuclear	14:16
Wind	18:34
Photovoltaic	3:6
Large hydropower	170:280
Plantation biomass	3:5

Source: *Comparing Energy Options with Life-cycle Assessment*, Luc Gagnon, Hydro Quebec, 2009.

energy payback ratios than shown in Table 2.4, which is based on silicon solar cells. Biomass scores similarly to coal- and gas-fired power plants on this measure because the power plant technologies are relatively similar. Hydropower, wind power, and nuclear power are the best-performing technologies based on these figures.

THE BOTTOM LINE

Most environmental assessments of power generation indicate that there are benefits to be gained in shifting from reliance on fossil fuels to other, primarily renewable, forms of generation. In most cases, however, the determining factor remains cost. Indeed cost has become more decisive over the last 20 years as the control over the power generation industry has shifted, in many parts of the world, from the public sector to the private sector.

The private sector requires short-term return on investment. This favors technologies that are cheap to build because loans for construction are small and can be repaid quickly. Most renewable technologies are capital intensive. The generating plant costs a lot to build but very little to run because the fuel (e.g., wind, sunlight, or water) is usually free. These plants are more cost effective over the long term, probably 20 years or longer, but less so over a shorter term.

Governments cannot direct the private sector but they can influence the industry with legislation, surcharges, and incentives. Such governmental tools are being used with some effect. Financial institutions are also beginning to heed the shift in consensus. In June 2003 a group of commercial banks agreed to a set of guidelines called the "Equator Principles," which are intended to provide a framework for assessing the social and environmental issues associated with a project seeking a loan. These guidelines are voluntary but potentially significant.

A shift away from fossil fuels will have a profound effect on the whole power generation industry. Not only generation but transmission and distribution management and structure will be affected. The change will, initially at least, be expensive. As a result, change will come slowly. What does appear clear in the second decade of the 21st century is that the change will come.

Chapter 3

Coal-fired Power Plants

Coal is the most important source of electrical power in the world today. At the end of the first decade of the 21st century it was responsible for over 40% of world electricity production, an annual output of around 8100 TWh out of a global total of 20,000 TWh in 2010, according to the U.S. Energy Information Administration.[1]

Significant deposits of coal can be found in most parts of the world, from the United States to South Africa, across Europe, many parts of Asia, and Australia. Exceptions exist, such as Japan and Taiwan, where resources are tiny; these countries, nevertheless, use large quantities of coal that they import. There are limited reserves in South America, and in Africa only South Africa has significant quantities while across the rest of that continent there is little.

The attractions of coal are that it is so widespread, it is often in plentiful supply, and where it is available it provides both a cost-effective and a secure source of electricity. It is for these reasons that many of the major economies across the world, including the two largest—the United States and China—have built their economic prosperity on coal.

Figures from the World Coal Association suggest that 45% of U.S. electricity is generated in coal-fired power plants, while in China the proportion is 79%. Table 3.1 lists the proportion of power derived from the fuel for these and the 10 other nations that rely most heavily on coal for electricity. The most reliant is South Africa, which derives 93% of all its electricity from coal, followed closely by Poland with 90%. India is high up the list with 69%, while Germany, last on the list, derives 44% of its electrical power from coal. Many of these countries have their own reserves but some, such as Morocco, import the fuel.

Alongside availability, the cost of coal is another important factor in its dominance. Coal prices have traditionally been stable and low compared to alternative fossil fuels, with coal typically half the cost of natural gas on an energy content basis. In recent years the price of the fuel has shown greater volatility than in the past, as has the price of many other globally traded commodities. Nevertheless, the cost is still low enough to make it an attractive source of electricity in many parts of the world.

1. *International Energy Outlook 2013*, U.S. Energy Information Administration.

TABLE 3.1 Coal-fired Power Generation for Leading Coal-using Countries*

Country	Proportion of Electricity Derived from Coal
South Africa	93%
Poland	90%
China	79%
Australia	76%
Kazakhstan	70%
India	69%
Israel	63%
Czech Republic	56%
Morocco	55%
Greece	55%
United States	45%
Germany	44%

**These data are for 2008–2009.*

Source: World Coal Association.

TABLE 3.2 Energy Densities for Different Fossil Fuels

Fossil Fuel	Average Energy Density
Coal	24 MJ/kg
Crude oil	46 MJ/kg
Natural gas	54 MJ/kg

On the other hand, coal has a lower energy density than other fossil fuels (Table 3.2) and is consequently more expensive to transport. When it is transported, it is usually by road over short distances; by rail, river, or canal over longer distances; and by ship when traded intercontinentally. Pipeline transportation, common with both oil and gas, cannot be used for coal. Therefore, coal is most economically utilized close to its source.

The major disadvantage of coal is that it is the dirtiest of the fossil fuels. Many coals contain significant amounts of sulfur, which when burned generates sulfur dioxide (SO_2). They can also contain numerous trace elements including heavy metals such as mercury, which are released when the fuel is burned and can find their way into the environment. Coal combustion, like all other types of fossil fuel combustion, generates nitrogen oxide (NO_x), which is damaging if released into the atmosphere, and the combustion process also releases large quantities of dust. On top of this, coal combustion generates more carbon dioxide (CO_2) for each unit of energy produced than any other fossil fuel.

In spite of all these potential pollutants associated with its combustion, for most of its history as a source of electricity energy, coal combustion has not been controlled. As a consequence it has been one of the most widespread and damaging forms of human industrial activity, generating high levels of pollution in many parts of the world. Coal combustion is now much more tightly regulated than in the past, and extensive emission control systems are required for new coal-fired power stations built virtually anywhere. This has made coal-fired power generation less cost effective than in the past, but it still remains competitive with other major sources of electrical energy.

TYPES OF COAL

The term *coal* embraces a range of materials. Within this range there are a number of distinct types of coal, each with different physical properties. These properties affect the suitability of the coal for power generation.

The hardest of coals is anthracite. This coal contains the highest percentage of carbon (up to 98%) and very little volatile matter or moisture. When burned it produces little ash and relatively low levels of pollution. Its energy density is generally higher than other coals at 23 MJ/kg to 33 MJ/kg. Anthracite is typically slow-burning and often difficult to fire in a power station boiler unless it is mixed with another fuel. While its energy content makes it attractive as a power plant fuel, the difficulty with firing it and its cost does not, so it has traditionally been used for heating rather than industrial use. However, it is becoming more common as a power plant fuel as countries with large reserves, such as Russia and Ukraine, switch to anthracite to free natural gas for export.

While anthracite reserves are important, the most abundant of the coals are the bituminous coals. These coals contain significant amounts of volatile matter. When they are heated they form a sticky mass, from which their name is derived. Bituminous coals normally contain between 45% and 70% carbon. Moisture content is between 5% and 10%. They burn easily, especially when ground or pulverized. This makes them ideal fuels for power stations. Bituminous coals are further characterized, depending on the amount of volatile matter they contain, as high-, medium-, or low-volatile bituminous coals. Some bituminous coals contain high levels of sulfur, which can be a handicap for power generation purposes.

A third category, called sub-bituminous coals or soft coals, are black or black-brown. These coals contain between 35% and 45% carbon and 15% to 30% water, even though they appear dry. They burn well, making them suitable as power plant fuels, and sulfur content is low.

The last group of coals that are widely used in power stations are lignites. These are brown rather than black and have a carbon content of 20–35%. Moisture content is 30–50%. Lignites are formed from plants that are rich in resins and contain a significant amount of volatile material. The amount of water in lignites, and their consequent low carbon content, makes the fuel uneconomic to transport over any great distance. Lignite-fired power stations are usually found adjacent to the source of fuel.

A type of unconsolidated lignite, usually found close to the surface of the Earth where it can be strip-mined, is sometimes called brown coal. (This name is common in Germany.) Brown coal has a moisture content around 45%. Peat is also burned in power plants, though rarely.

COAL RESERVES

Table 3.3 shows the global proven coal reserves, which at the end of 2010 amounted to 860,938 Mtonnes. The greatest reserves are in Europe and Eurasia with 304,604 Mtonnes, followed by the Asia-Pacific region with 265,843 Mtonnes and North America with 245,088 Mtonnes. Reserves in the Middle East and Africa are small by comparison at 32,895 Mtonnes, and those in Central and South America even smaller at 12,508 Mtonnes.

When these are subdivided into types as shown in the table, Europe's reserves can be seen to be composed primarily of lower-quality sub-bituminous coals and lignite. North America has relatively higher quantities of the high-quality coals but the greatest reserves of these are in the Asia-Pacific region.

TABLE 3.3 Proven Coal Reserves at the End of 2010

Region	Anthracite and Bituminous (Mtonnes)	Sub-bituminous and Lignites (Mtonnes)	Total (Mtonnes)	Reserve/ Production Ratio
North America	112,835	132,253	245,088	228
Central and South America	6890	5618	12,508	124
Europe and Eurasia	92,990	211,614	304,604	242
Middle East and Africa	32,721	174	32,895	126
Asia-Pacific	159,329	106,512	265,843	53
Total	404,762	456,176	860,938	112

Source: BP Statistical Review of World Energy.

While coal is generally considered to be widely available, five countries hold 75% of all proven coal reserves: United States USA (28%), Russia (18%), China (13%), Australia (9%), and India (7%). At the current rate of consumption the world's proven coal reserves would last for another 112 years, as shown in Table 3.3. However, this varies from region to region with European and Eurasian reserves likely to persist the longest, while those in the Asia-Pacific region might only last for a further 57 years if current production rates continue. However, in all regions proven reserves represent only a portion of the coal that actually lies in the ground, and these potential reserves are likely to be much higher than those shown in Table 3.3.

The leading coal-producing nation is China, which mined 3520 Mtonnes of hard coal in 2011. The United States, the second largest producer, mined only 993 Mtonnes, and India 589 Mtonnes. Other important coal-producing countries include Australia (416 Mtonnes in 2011), Indonesia (325 Mtonnes), South Africa (255 Mtonnes), and Russia (334 Mtonnes). The main European coal producers are Germany (189 Mtonnes) and Poland (139 Mtonnes). According to the World Coal Association most coal is consumed in the country of origin and only around 16% is traded internationally. However, this represents an important source of energy for countries, such as Japan, South Korea, and Taiwan, which have virtually no reserves of their own. China, in spite of its massive coal production, is also a net importer of coal. Meanwhile, the major exporting nations are Australia, Indonesia, and Russia.

The cost of transporting coal means that trade routes are normally kept as short as possible. So, for example, Australian and Indonesian coal is most often sold to other Asia-Pacific countries, such as Japan, South Korea, and China. Russia and South Africa, another important coal-exporting nation, trade mostly with Europe, while U.S. coal exports to other countries in the Americas.

COAL CLEANING AND PROCESSING

Coal cleaning offers a way of improving the quality of a coal, both economically and environmentally. The most well-established methods of coal cleaning focus on removing excess moisture from the coal and reducing the amount of incombustible material, which will remain as ash after combustion. Moisture removal reduces the weight and volume of the coal, rendering it more economical to transport and easier to burn. Ash removal improves its combustion properties and aids power plant performance.

Moisture is removed from coal by drying. This can simply be solar drying—leaving the coal in the open before transporting it. Drying coal in this way can reduce its mass and increase its energy density, making it relatively cheaper to transport.

The alternative—drying coal by heating—is most often carried out at the power station, utilizing waste heat in the plant flue gases. Such a procedure is absolutely essential when burning high-moisture lignites such as brown coal.

In this case, drying does not affect transportation costs because the fuel has, by this stage, already reached the power station.

Ash removal is carried out by crushing the coal into small particles. Incombustible mineral particles are more dense than the coal and can be separated using a gravity-based method. Such treatment will remove some minerals containing sulfur, and can result in a reduction of up to 40% in sulfur dioxide emissions during combustion. (Some sulfur is bound to the carbon in the coal; such sulfur is not affected by this type of cleaning.)

There have been attempts to develop more advanced methods for coal treatment employing either higher-temperature processing of the coal or chemical rather than physical processes. These processes, which are aimed at removing polluting impurities in the coal to make it a cleaner fuel to burn, have not so far found commercial application. Coal-cleaning processes are also being developed to treat coal wastes that have previously been discarded to make them suitable for combustion.

TRADITIONAL COAL-FIRED POWER GENERATION TECHNOLOGY

The traditional method of producing electricity from coal is to burn the fuel in air and capture the heat released during the combustion in a boiler where it is used to raise steam, the latter driving a steam turbine generator. This type of power plant has been evolving since the modern steam turbine was invented by Charles Parsons in 1884. His first unit drove a dynamo that produced 7.5 kW of electrical power.

A power plant of this type is made up of several key components (Figure 3.1). There will be a fuel handling system designed to receive the bulk

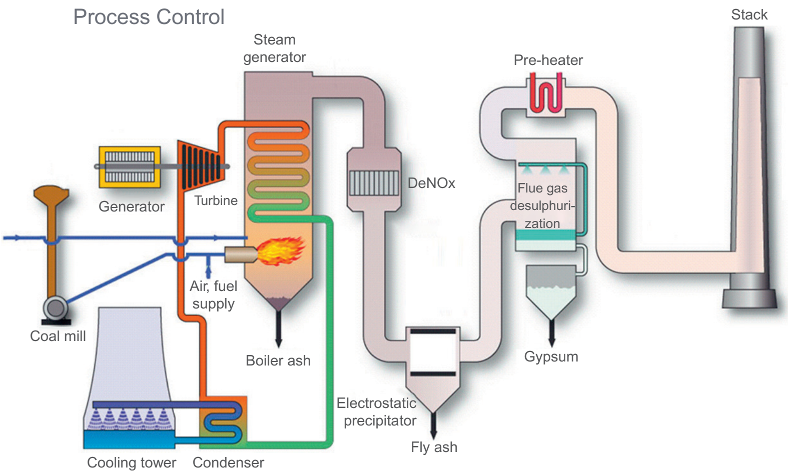

FIGURE 3.1 Schematic of a pulverized coal-fired power station.

coal deliveries from the mine and convert the mined coal into a form that can be readily burned. In most modern coal-fired power plants this involves crushing the fuel to produce a powder. This pulverized coal is then fed into a combustion system where it is mixed with air and ignited under controlled conditions, releasing chemical energy as heat. This heat is captured by water within tubes in the boiler, the heat energy converting the water into steam. The combustion system and boiler must be closely integrated for highest efficiency and will normally be considered a single unit in modern plants. Once combustion is complete most of the ash residue falls to the bottom of the combustion chamber and is removed as slag. However, some ash forms into fine particles that are carried away with the hot combustion gases. These particles must be removed at a later stage.

Steam generated in the boiler is carried to a steam turbine that is designed to extract as much heat energy as possible from the gas. In a large coal-fired power station the steam turbine is likely to be composed of at least three elements: a high-pressure (HP) turbine, an intermediate-pressure (IP) turbine, and one or more low-pressure (LP) turbines. Steam passes from one to the next in sequence. Higher efficiency can often be achieved by reheating the steam between the HP and IP turbines by returning it through a special stage of the boiler called a reheater.

To extract the maximum amount of energy from the steam, a condenser is fitted to the output of the LP turbine(s) to condense the steam back into water. The colder the cooling water used in the condenser, the higher the efficiency will be. The water is then returned to the boiler and passes around the cycle again.

The hot flue gases that exit the combustion chamber of the plant and pass through the boiler are rich in carbon dioxide and laden with impurities. These impurities include sulfur dioxide, nitrogen oxide, heavy metals, organic compounds, and tiny ash particles. All these impurities must be removed before the flue gas can be released into the atmosphere. Cleaning is carried out in a sequence of flue-gas cleaning systems, each a separate chemical or filtration plant. Carbon dioxide in not currently removed from commercial coal-fired power plants, but technology for its capture is advanced and is expected to be deployed within a decade. A plant that applies all these processes, including carbon dioxide capture, may be called a zero-emission plant, although in fact, traces of all will still be released.

Efficiency is the key to modern coal-burning technology. The higher the ratio of electrical energy output to chemical energy input of the coal combustion process, the cheaper each unit of electricity produced will be. For modern plants without carbon dioxide capture, higher efficiency also means lower emissions per unit of electricity produced.

Of the chemical energy contained with the coal, around 15% is lost to the energy conversion system. The remainder is utilized to heat steam so that the hot steam contains around 85% of the original chemical energy. Converting the hot steam into electricity relies on the Carnot thermodynamic cycle. Conversion efficiency depends on the temperature and pressure of the steam (more accurately the temperature and pressure drop that is achieved between steam turbine inlet and outlet), so the development of coal-fired power plant technology is directed at producing steam at the highest temperature and pressure possible. From an energy viewpoint, therefore, the two most important components of a coal-fired power station are the boiler, which produces high-temperature, high-pressure steam, and the steam turbine, which must then convert the energy carried by that steam into electrical energy.

BOILER TECHNOLOGY

The boiler in a coal-fired power plant converts chemical energy contained within the coal into heat energy that is captured and carried away in hot, high-pressure steam (Figure 3.2). Coal-burning technology has a long history and has seen many plant configurations, but the most important type of plant in modern use is the pulverized coal-fired power plant, often abbreviated as PC plant. PC plants probably account for more than 90% of all operating coal-fired power stations in general use and this type of plant is the main plant design for future high-efficiency coal-fired plants.

A PC plant burns coal that has first been ground to a fine powder using large grinding mills. A typical plant will have several of these, each feeding a single

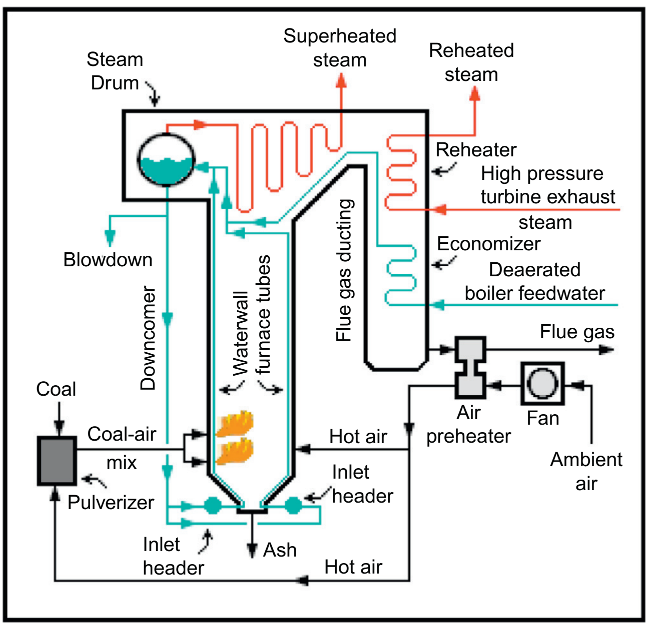

FIGURE 3.2 Cross-section of a pulverized coal plant boiler.

burner. These burners are where the coal, mixed with air, is injected into the boiler where it ignites in a high-temperature fireball inside the furnace, consuming the fuel and releasing chemical energy as heat.[2] Several burners are used to create a stable fireball in the center of the furnace where combustion takes place. The temperature within the fireball may reach 1500–1700 °C in the hottest part of the flame. At these temperatures the nitrogen in air is easily oxidized to produce nitrogen oxide. To keep nitrogen oxide production at a minimum, the amount of air entering the combustion chamber with the pulverized coal is restricted, maintaining a chemically reducing atmosphere in which virtually all the oxygen is captured during carbon combustion, leaving none to react with nitrogen. Additional air is then introduced higher up the combustion chamber to complete the coal combustion process, at a point where the temperature of the combustion gases has lowered and the reaction of oxygen with nitrogen proceeds more slowly.

The heat released during combustion is partly radiant heat and partly convective heat. Radiant heat is captured by water running in tubes within the walls of the combustion chamber. Further collections of tubes are placed in the path of the flue gases exiting the combustion chamber, and as water passes from one set of tubes to another its temperature rises and finally steam is generated (Figure 3.3). In a conventional boiler this will take place within a steam drum, which allows the phase change between liquid and gas to proceed smoothly. Steam from the drum may then be superheated to create an even higher-temperature gas.

In the most modern plants the temperature and pressure within the boiler tubes is such that water enters a supercritical state under which the distinction between liquid and gaseous states cease to exist. With this type of plant a steam drum is not necessary. Units of this sort are called supercritical boilers, while those with steam drums are called subcritical boilers. Supercritical boilers are capable of producing steam at much higher temperatures and pressures than subcritical boilers, therefore steam turbines fed from such plants are capable of higher Carnot cycle efficiency.

The supercritical point for water occurs at 22.1 MPa/374.1 °C. So long water remains above the supercritical point and within its supercritical phase there is no longer a need for a drum in a power plant steam system to allow the water to boil. Instead the conversion from water to steam can occur homogenously within the boiler pipes as the temperature rises.

Steam conditions at the exit of a modern subcritical power plant boiler vary widely but typical figures would be a steam temperature of 540 °C and a steam pressure of 17.5 MPa. Such a plant would be capable of generating power at an efficiency of around 38%, depending on site conditions. Supercritical power plants operate at steam exit temperatures of 540–600 °C and at exit pressures of 23–30 MPa. These plants can achieve efficiencies up to 41%. There is a

2. There is normally a facility within the plant to use to inject light oil or gas into the plant to initiate combustion within the furnace.

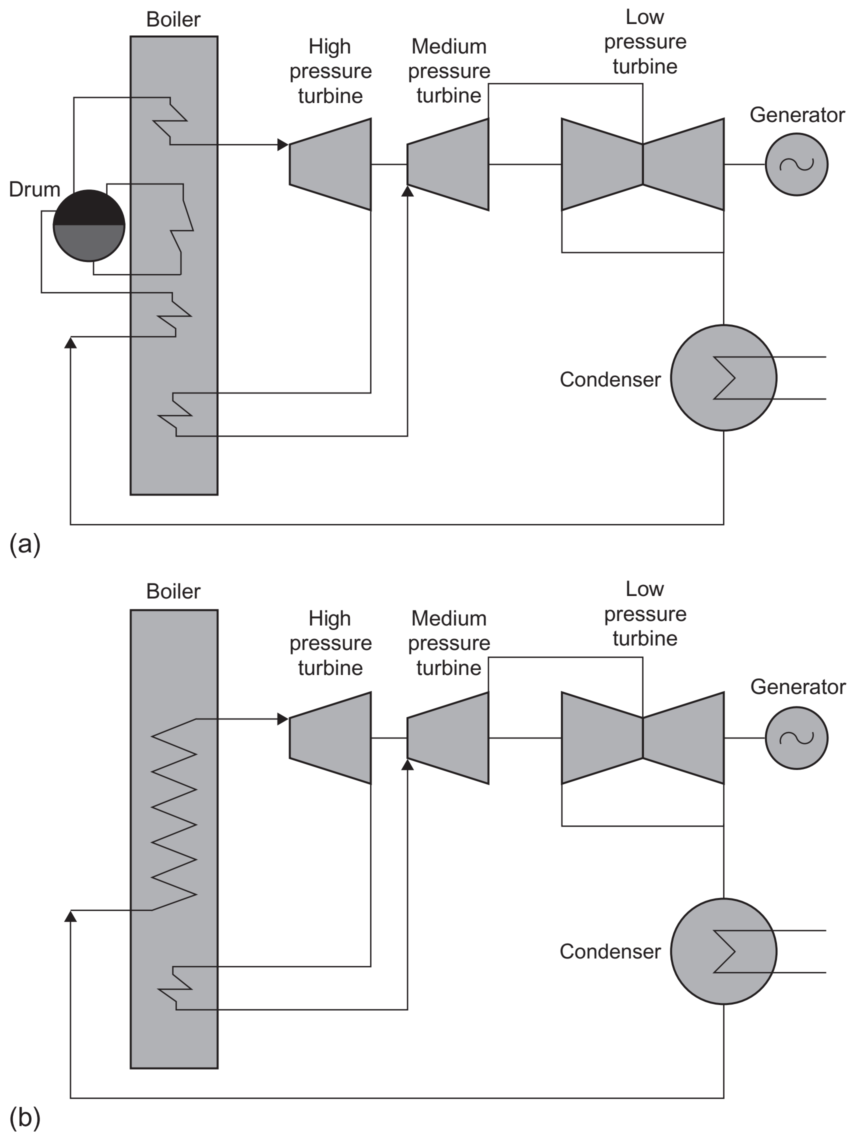

FIGURE 3.3 Coal-fired power station boiler steam cycles: (a) typical subcritical steam cycle with a conventional drum boiler and natural circulation and (b) typical supercritical steam cycle with once-through boiler.

further subdivision called ultra-supercritical power plants that operate at even higher temperatures and pressures. Although the definition of ultra-supercritical is not precise,[3] plants operating at these more extreme conditions have been

3. The U.S. Electric Power Research Unit has defined an ultra-supercritical power plant as one operating with a boiler exit steam temperature above 593 °C.

capable of up to 45% efficiency under optimum site conditions. In comparison, the average efficiency of coal-fired plants operating across the globe is around 28% and that of the U.S. coal fleet is around 33%.

The extreme temperatures and pressures in supercritical boilers make extreme demands on the materials used to construct them. While steam boilers have traditionally been constructed from steels, conditions in ultra-supercritical power plants require that some components are made from nickel-based alloys similar to those used to construct the high-temperature components used in gas turbines. These materials are more costly than steels. With further materials development it is expected that future plants will be able to operate at steam temperatures of 700–750 °C. This should permit an energy-to-electrical conversion efficiency as high as 55% to be achieved in a coal-fired steam plant.

While efficiency is the most important factor driving boiler design, flexibility has also been recognized as vital in recent years. Coal-fired power plants have traditionally operated as base-load power stations operating essentially at full output all the time. This is no longer the situation everywhere. In some regions coal-fired power stations are being used to support the generation of renewable electricity. This means they have to be able to operate both efficiently and effectively at part load as well as full load, and to be able to change output load as required by the grid. One technique being used to achieve this is sliding-pressure operation under which the steam pressure is allowed to fall as output falls but steam temperature is maintained. With sliding-load operation it is possible to maintain relatively high efficiency at part load, even though this may involve falling below the critical point of water.

STEAM TURBINE DESIGN

The steam turbine is the primary mechanical device (sometimes called the prime mover) in most conventional coal-fired power stations. Its job is to convert the heat energy contained in the steam exiting the boiler into mechanical energy, rotary motion. The steam turbine first appeared in power applications at the end of the 19th century. Before that steam power was derived from steam-driven piston engines.

The steam turbine is something of a cross between a hydropower turbine and a windmill (Figures 3.4 and 3.5). It, like them, is designed to extract energy from a moving fluid and the fluid it exploits is a form of water, the same as the hydroturbine. In the case of a hydroturbine the water remains in the liquid phase and neither its volume nor its temperature changes during energy extraction. In the case of the steam turbine, energy extraction is from a gas (steam) rather than a liquid and involves both the pressure and temperature of the fluid falling. This has a profound effect on the turbine design.

Both hydroturbines and steam turbines exist in two broad types: impulse turbines that extract the energy from a fast-moving jet of fluid, and reaction turbines that are designed to exploit the pressure of a fluid rather than its motion.

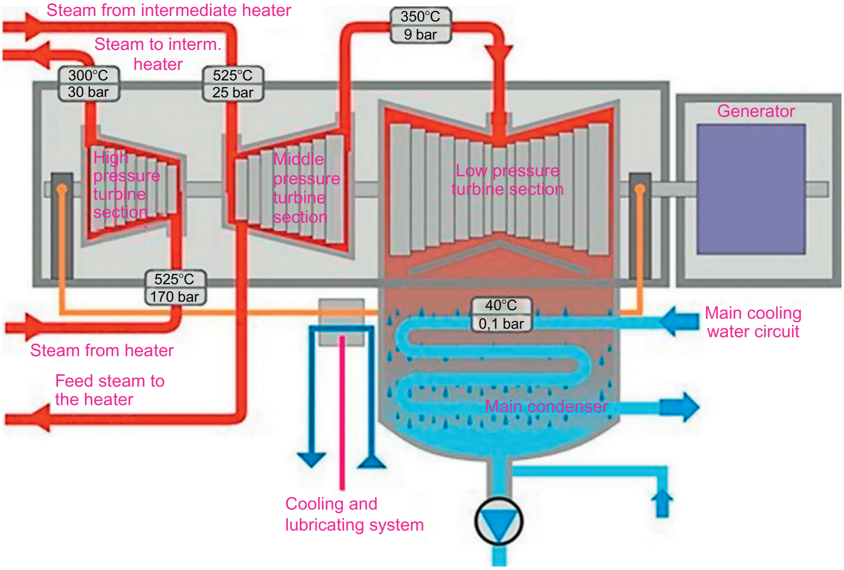

FIGURE 3.4 Schematic of a large steam turbine for a coal-fired power plant.

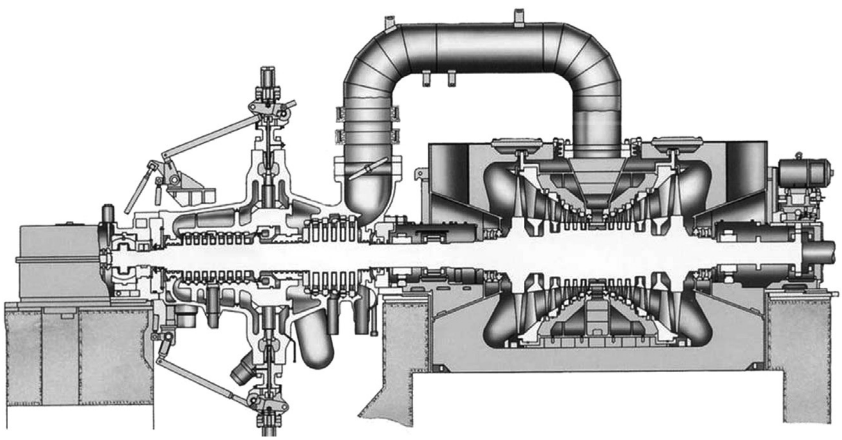

FIGURE 3.5 Section through a modern steam turbine. *Source: Toshiba Industrial and Power Systems & Services Company.*

A hydroturbine will be of one design or the other. In a steam turbine the two principles may be mixed in a single machine and they may even be mixed in a single turbine blade.

It is impossible to extract all the energy from steam using a turbine with a single set of turbine blades. Instead, a steam turbine utilizes a series of sets of blades, called stages. Each stage is followed by a set of stationary blades (usually called nozzles) that control the steam flow to the next stage.

A single steam turbine stage consists of a set of narrow blades projecting from a central hub. (In concept, it is something like a steam windmill.) Ten or more sets of blades can be mounted on a single steam turbine shaft. This combination of shaft and blades is called a rotor. The turbine stages are separated by carefully designed stationary blades, or nozzles, that control the flow of steam from one set of rotating blades to the next. The precise shape of the blades in each set determines whether that set is impulse or reaction, or a cross between the two. The hub, blades, and nozzles are enclosed in a close-fitting case to maintain the steam pressure.

In a steam turbine impulse stage, energy is extracted at constant pressure while the velocity of the steam falls as it flows across the blades. The steam is then expanded through a stationary control stage to increase its velocity again before energy is extracted from another set of impulse blades. In a steam turbine reaction stage, by contrast, both pressure and velocity of the steam fall as energy is extracted by the rotating blades.

Steam exiting the power plant boiler is at a high temperature and pressure. Both temperature and pressure fall as the steam passes through the turbine. The greater the temperature drop and the greater the pressure drop available, the more energy can potentially be captured from the steam. Consequently, the most efficient power plants condense the steam back to water at the end of the turbine.

Even with a modern design it is impossible to capture all the energy from the steam efficiently with a single multiple-stage turbine. Coal-fired power plants use several, designated HP, IP, and LP turbines. The blades in these turbines get larger (longer) as the pressure drops; in fact, the LP turbine may comprise several turbines operating in parallel to extract the last energy from the steam because a single turbine designed to achieve the same energy extraction would be impossibly large. All the turbines may be mounted on a single shaft, but it is common for the LP turbines to be on a separate shaft rotating at a lower speed to reduce the forces exerted as the blade tips. Multiple turbines of this type can have aggregate outputs over 1000 MW.

As with boilers, the demands of modern power plant design have led to the development and introduction of high-performance materials that can cope with the extreme conditions encountered within a steam turbine. The high-pressure turbine blades have to be able to withstand extremes of both temperature and pressure and to resist the abrasive force of steam. At the low-pressure end of the turbine train the large size of the turbines means that the blade tip speeds are enormous, again requiring specially designed materials to withstand the centrifugal forces exerted on them.

A refinement that improves the overall efficiency in a steam plant is to return the steam to the boiler after it has passed through the HP turbine, reheating it before delivering it to the IP turbine. Most modern steam turbine plants use this single reheat design (multiple reheat is also possible).

The theoretical maximum efficiency of a coal-fired power station is determined by the temperature difference between the steam entering the HP turbine and the temperature exiting the LP turbine. The greater this temperature difference, the more energy can be extracted. With the most advanced technology, utilizing the best boiler materials to achieve the highest steam temperatures and pressures and with optimum site conditions, a maximum efficiency of 43–45% can be achieved. New supercritical designs may eventually push this as high as 55%. In the near future, however, the best that is likely to be achieved is something between 47% and 49%.

GENERATORS

The turbine shaft, or shafts if there are more than one, are coupled to a generator that converts the rotary mechanical motion into the electrical energy that the plant is designed to produce. Generators, like steam turbines, first appeared during the 19th century. All utilize a coil of a conducting material, usually copper, moving in a magnetic field to generate electricity, as shown in Figure 3.6.

The generators used in most power stations, including coal-fired power stations, are designed to deliver an alternating current (AC) to a power grid. An AC current is preferred because it allows the voltage to be raised or lowered easily using a transformer. For transmission of power over long distances it is preferable to use a very high voltage and a low current. The voltage is then reduced with a transformer before delivery to the consumer.

The need to generate an AC voltage determines the speed at which the generator rotates. This must be an exact multiple of the grid frequency (normally grids operate at either 50 Hz or 60 Hz). For grids operating at 50 Hz the traditional generator speed is 50 cycles per second, or 3000 rpm for a standard two-pole generator, or half this speed, 1500 rpm, for a four-pole generator. The equivalent 60 Hz two-pole machine rotates at 3600 rpm and a four-pole machine at 1800 rpm. This speed, in turn, determines the operating speed of the steam turbine. Large LP steam turbines may operate at the lower speeds to reduce stress on the long turbine blades.

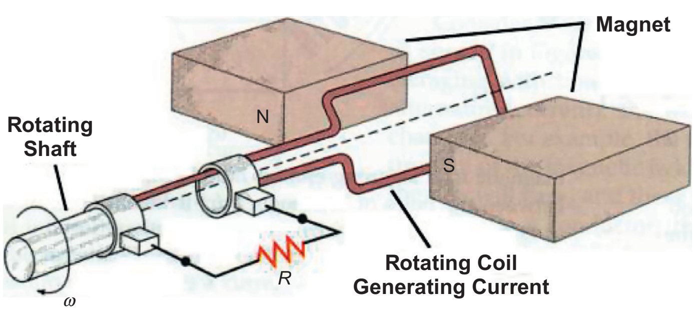

FIGURE 3.6 Simple schematic of a power plant generator.

Generators may be as large as 2000 MW, and large generators are normally designed for the specific project in which they will operate. Modern generators operate with an efficiency of greater than 95%. The remaining 5% of the mechanical input energy from the turbine is usually lost as heat within the generator windings and magnetic components. Even though the percentage is small, this still represents an enormous amount of energy, perhaps 50 MW in a 1000 MW machine. Therefore, generators require very efficient cooling systems to prevent them from overheating. For smaller generators up to around 300 MW, air cooling will normally be sufficient. Larger generators, up to 500 MW, usually use hydrogen cooling. Above this size, the rotating section of the generator, the rotor, will normally be hydrogen cooled while the stationary part, the stator, will be water cooled through pipes embedded into its structure.

The broad outline of generator design has changed little over a century. However, new materials have improved efficiencies. The latest developments involve the use of superconducting materials to reduce energy loss and increase efficiencies. These have yet to achieve commercial viability.

FLUIDIZED BED COMBUSTION

While the pulverized coal-fired boiler represents the most efficient type of coal-fired power plant, there are alternatives. One of these is the fluidized bed boiler. This type of combustion system is capable of burning a much wider range of fuels than a PC plant. It operates at a lower temperature, making the production of nitrogen oxide less of a problem, and it can also incorporate sulfur capture within the combustion chamber, leading in principle at least to easier emission control.

The concept behind the fluidized bed is simple: to create a solid-state reactor that mimics liquid-phase reactors. If a layer of sand, finely ground coal, or another fine solid material is placed in a container and high-pressure air is blown through it from below, the particles, provided they are small enough, become entrained in the air and form a floating (i.e., fluidized) bed of solid particles above the bottom of the container. This bed of solid particles now behaves like a liquid in which the constituent particles constantly move to and from and collide with one another like the molecules in a liquid. As a type of reactor, this offers some significant advantages because it encourages much more rapid and thorough reaction between the particles within the bed or between the particles and air in the case of combustion.

The fluidized bed was used first in the process industries to enhance the efficiency of chemical reactions between solids by simulating conditions of a liquid-phase reaction and as a simple method of reducing emissions from industrial plants. Only later was its application for power generation recognized. Its use is now widespread. The fluidized bed can burn a variety of coals and other poorer fuels, such as peat, coal-cleaning waste, petroleum coke, wood, and other biomass. See Figure 3.7 for a diagram of its circulation.

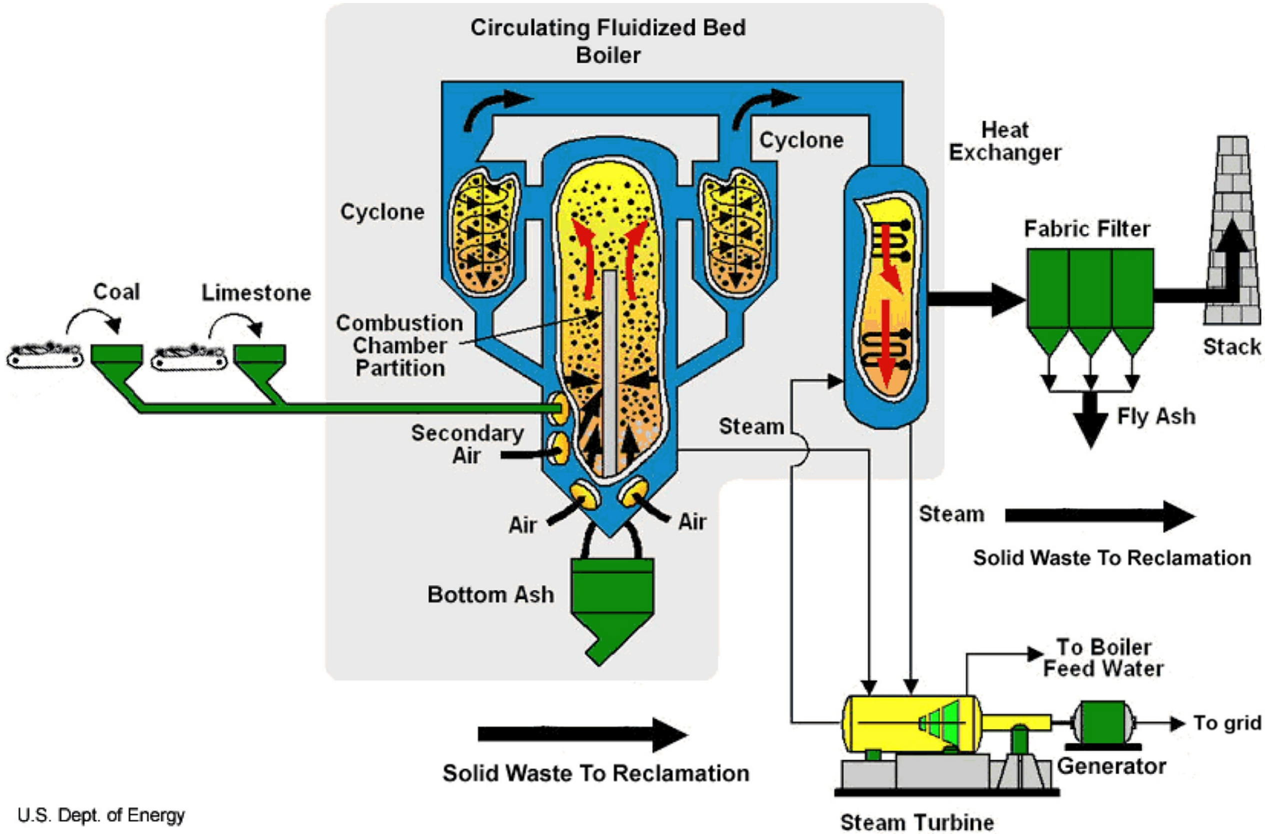

FIGURE 3.7 Circulation of a fluidized bed boiler.

A fluidized bed used for power generation contains only around 5% coal or fuel within the actual bed. The remainder of the bed is primarily an inert material, such as ash or sand. The temperature in a fluidized bed is around 950 °C, significantly lower than the temperature in the heart of a pulverized coal furnace. This low temperature helps minimize the production of nitrogen oxide. A reactant such as limestone can also be added to the bed to capture sulfur, reducing the amount of sulfur dioxide released into the exhaust gas. One further advantage of the fluidized bed is that boiler pipes can be immersed in the bed itself, allowing extremely efficient heat capture (but also exposing the pipes to potentially high levels of erosion).

There are three primary designs for fluidized bed power plants in use. The simplest, called a bubbling bed plant, is essentially a conventional boiler in which the combustion chamber has been replaced with a fluidized bed. Air is blown at relatively low velocity into the chamber from beneath the bed to maintain its fluid state and fuel introduced from above. Ash is removed from below, as in a PC plant, and additional, overfire air is blown in above the bed to complete the combustion process. Bubbling fluidized bed plants are often used in biomass plants.

The second, more complex design is the circulating fluidized bed plant. In this type of fluidized bed the particles are fluidized at high speed by using high-velocity air and a cycling fluidized mass is created. The high-speed particles within this type of plant pass up out of the combustion chamber and are then recovered and recirculated by passing the flow of flue gases and particles through a cyclone filter that returns the particles to the bed while allowing the flue gases to flow onwards. Such plants provide good mixing and long residence times for particles so that emission control using additives in the bed can be more effective. They are also capable of burning higher calorific value fuels such as anthracite, which would be more difficult to combust completely in the bubbling bed reactor.

The circulating bed can remove 90–95% of the sulfur from the coal while the bubbling bed can achieve between 70% and 90% removal. Maximum energy conversion efficiency is 43%, similar to that of a traditional pulverized coal plant. However, such high efficiencies can only be achieved with larger plants that can employ larger, and generally more efficient, steam turbines under optimum steam conditions. Advanced fluidized bed plants are built with supercritical boilers to achieve these high conversion efficiencies.

A third type of fluidized bed design, called the pressurized fluidized bed (Figure 3.8), was developed in the late 1980s and the first demonstration plants employing this technology were constructed in the mid-1990s. The pressurized bed is like a bubbling bed, but operates at a pressure between 5 atmospheres and 20 atmospheres. This allows higher overall efficiency to be achieved.

Operation under high pressure means that the flue gases exiting the boiler can be passed through a gas turbine to capture additional energy for power generation. This is in addition to the power generated by the traditional steam cycle.

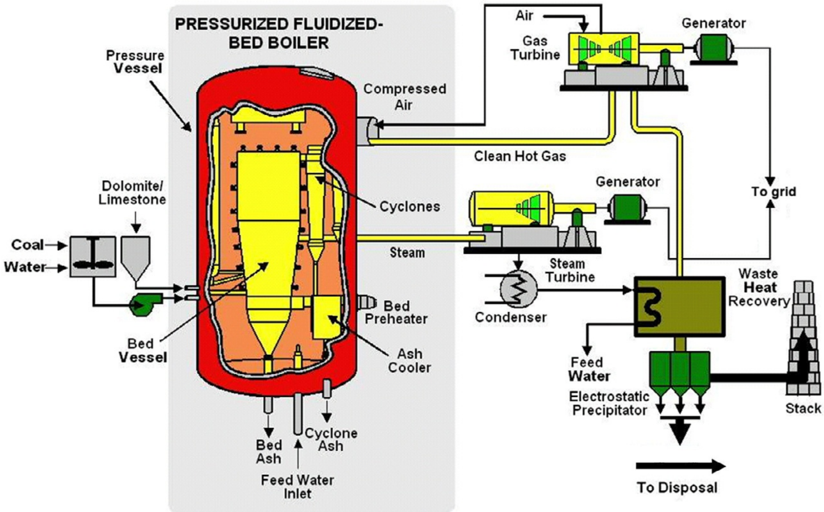

FIGURE 3.8 Pressurized fluidized bed combustion plant schematic.

In principle, a pressurized fluidized bed power plant is capable of a conversion efficiency of 44–45%, but in practice they have not achieved this level. The largest pressurized fluidized bed power plant in operation is a 360 MW unit in Japan. This plant, with a supercritical boiler, has a claimed efficiency of 42%.

Although commercial fluidized bed power plants with capacities of over 400 MW are available, the units have not proved as efficient as supercritical PC power plants for coal combustion. Their sulfur and nitrogen oxide emissions, though lower than from an unmitigated PC power plant, are often too high to meet statutory regulations and additional measures are required to control the emissions. However, the plants can burn a variety of fuels other than coal. This has made them attractive for a range of projects where low-grade fuels, fuel waste, or biomass are to be burned.

INTEGRATED GASIFICATION COMBINED CYCLE

A second alternative to the pulverized coal-fired power plant is the integrated gasification combined cycle (IGCC) plant that is based around the gasification of coal (Figure 3.9). Coal gasification is an old technology. It was widely used to produce town gas for industrial and domestic use in the United States and Europe until natural gas became readily available.

Modern gasifiers convert coal into a mixture of hydrogen (H) and carbon monoxide (CO), both of which are combustible. Gasification normally takes place by heating the coal with a mixture of steam and oxygen (or, in some cases, air). This can be carried out in a variety of gasifier designs including fixed bed, fluidized bed,

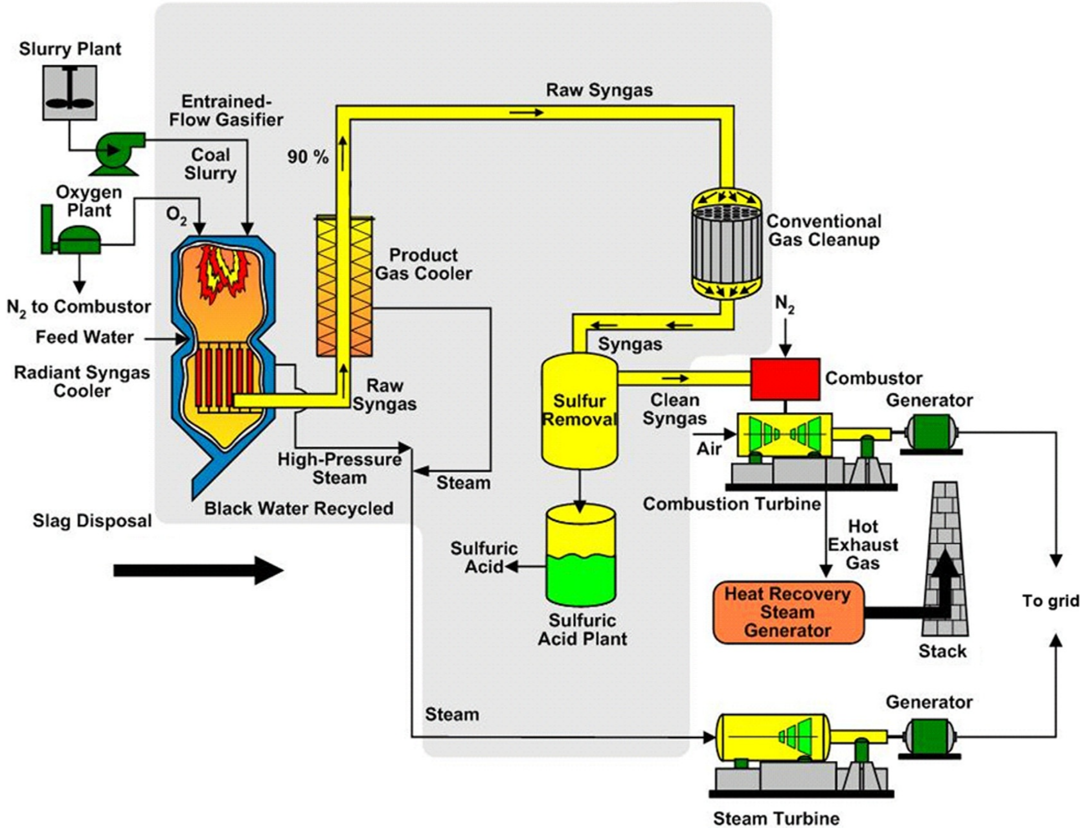

FIGURE 3.9 Flow diagram of an IGCC plant. *Source: Tampa Bay Electric Company.*

and entrained flow gasifiers. The primary reaction for combustible gas production is that of the carbon in coal with water to produce hydrogen and carbon monoxide:

$$C + H_2O = H_2 + CO$$

However, there are other reactions that produce carbon dioxide and methane or other hydrocarbons. Both hydrogen and carbon monoxide are combustible.

The process that takes place in the gasifier is a partial combustion of the coal. Consequently, it generates a considerable amount of heat. This heat can be used to generate steam to drive a steam turbine. The gas produced, meanwhile, is cleaned and can be burned in a gas turbine to produce further electricity. Heat from the exhaust of the gas turbine is used to raise additional steam for power generation. This is the basis of the IGCC plant although a variety of configurations exist.

There are a limited number of IGCC plants in operation with capacities of 250–300 MW. They have demonstrated efficiencies of 40%. However, capital costs are higher than for pulverized coal-fired power plants. Costs might improve if hot-gas cleanup technologies can be developed, but current IGCC technology does not offer a viable alternative to a supercritical PC plant under most conditions. However, IGCC might offer an effective means of carbon capture in a coal-fired power plant (see the following section) and this may prove its greatest opportunity for the future.

EMISSION CONTROL FOR COAL-FIRED POWER PLANTS

The combustion of coal is the dirtiest of the large-scale methods of generating electricity, primarily because of the range of pollutants that are found within the fuel. Most coals contain some sulfur. Often it is more than 3% of the coal and it may reach as much as 10%. When the coal is burned this sulfur is converted into sulfur dioxide and carried off by the flue gases. If released into the atmosphere it can be converted into an acid with potent consequences. There is a small amount of organic nitrogen with coal too. During combustion this is converted into nitrogen oxides of various sorts. However, the main source of these gaseous nitrogen compounds is the nitrogen in air that can become oxidized at the high temperatures encountered within coal furnaces. These two can be potent pollutants.

Coal usually contains a significant amount of mineral impurity too. A large part of this fuses to create solid lumps, which are left behind in the combustion chamber as slag. However, some are reduced to tiny solid particles that get carried away with the flue gas. The particles may contain heavy metals, such as cadmium and mercury, that, if allowed to escape, will be released into the environment. Some coals, particularly the bituminous varieties, contain large amounts of volatile organic compounds and these, or fragments of them generated by their incomplete combustion, can also be released. Incomplete combustion of the carbon in coal may also lead to significant levels of carbon monoxide within the flue gases.

Environmental regulations require that as far as possible these materials are removed from coal-fired power plant flue gases before the latter are released into the atmosphere. Different techniques have been developed for the most important of these: sulfur scrubbers for removing sulfur compounds, low nitrogen oxide burners and catalytic reduction systems to remove nitrogen oxide, and filters and electrostatic precipitators to control dust emissions. Other trace elements such as heavy metals may require their own removal plants but often these can be tackled alongside one or the other pollutants, making an additional chemical treatment process unnecessary.

Table 3.4 contains figures for the concentrations of various power plant airborne pollutants that are considered permissible in the EU and United States if good air quality is to be maintained. EU regulations are generally stricter; for example, the EU expects sulfur dioxide concentrations over a 24-hour period to be below 125 μg/m^3. In the United States the same standard is 365 μg/m^3. However, internationally, standards are tending to converge as the effects of even low levels of pollution on human health become more widely recognized. The PM10 particulate matter standard is for dust particles greater than 10 μm in diameter; this is generally the standard of importance when considering dust from coal-fired power plants. There are other standards including PM2.5 for particles up to 2.5 μm in diameter.

The figures in Table 3.4 apply to the air quality that people will encounter in the street or in their houses or offices when carrying out their daily lives. The actual emissions permitted by power plants are generally much higher than this. A power plant represents a concentrated source of pollutants, but these are released in hot gases from a tall stack so that they should rise high into the atmosphere and become diluted before humans or other life-forms come into contact with them. However, the behavior of the pollutants once they enter the atmosphere is not always predictable. The behavior of the plume of exhaust gases from a power plant stack will depend on atmospheric conditions, so sometimes the pollutants will fall close around the plant and at other times they may be carried across continents.

Table 3.5 shows some of the emission levels permitted within the EU for power plant flue gases. The figures are for large plants with a thermal capacity

TABLE 3.4 Air-quality Standards

Pollutant	EU Standard	U.S. Standard
Sulfur dioxide	125 μg/m^3	365 μg/m^3
Nitrogen oxide	40 μg/m^3	100 μg/m^3
Particulate matter (PM10)	40 μg/m^3	150 μg/m^3
Carbon monoxide	10 mg/m^3	10 mg/m^3
Ozone	120 μg/m^3	150 μg/m^3

Source: EU Commission and U.S. Environmental Protection Agency.

TABLE 3.5 EU Emission Limits for Large Power Plants

Sulfur dioxide emissions for plants built after 2003	200 mg/m^3
Sulfur dioxide emission limits after 2016	150 mg/m^3
Nitrogen oxide emissions for plants built after 2003	200 mg/m^3
Nitrogen oxide emission limits after 2016	150 mg/m^3
Dust emission limits after 2016	20 mg/m^3
Proposed mercury emission limit	30 μg/m^3

Source: EU Commission.

in excess of 300 MW. The limits are less strict for some smaller plants. For sulfur dioxide the limit for plants built after 2003 is 200 mg/m^3, falling to 150 mg/m^3 after 2016. Permitted emission levels for nitrogen oxide are the same. Dust emissions are to be below 20 mg/m^3 after 2016 and there is a proposed emission limit for mercury of 30 μg/m^3. As earlier, these EU limits are probably some of the strictest to be found, but as with air-quality standards, the regulations are becoming stricter everywhere.

There is one other important combustion product of coal combustion not included in the preceding tables or discussion: carbon dioxide. This is the reaction product when carbon is burned in air, the reaction of which releases the heat energy used to generate electricity. The flue gases from the boiler of a typical advanced coal-fired power plant may contain up to 14% carbon dioxide depending on the specific plant conditions.

The release of carbon dioxide from the combustion of fossil fuels in power plants and elsewhere into the atmosphere is widely regarded as the main cause for a steady but accelerating rise in average global temperatures over the past 150 years. This is seen as potentially damaging to the global environment. The capture and removal of carbon dioxide from fossil fuel power plant flue gases is not yet mandatory anywhere, but measures to try and control its emissions are being introduced in some parts of the world, particularly the EU. At the same time, methods for capturing the gas are being developed and there is a growing consensus that these will need to be deployed on a commercial scale after 2020 if global warming is to be limited. If this becomes necessary, then coal-fired power plants will be in the front-line since they are the greatest emitters.

COAL TREATMENT

Cleaning coal prior to combustion can significantly reduce the levels of sulfur emissions from a power plant, as well as reduce the amount of ash and slag produced. Physical cleaning, as outlined earlier, can have a beneficial effect on plant performance and maintenance schedules. It has been estimated that boiler availability improves by 1% for every 1% reduction in ash content. This may be cost effective, depending on the quality of coal being treated. However, more complex cleaning procedures have not so far proved economical.

One disadvantage of coal cleaning is that it leads to a loss of between 2% and 15% of the coal with the coal waste. It may be possible to burn this coal waste in a fluidized bed combustion plant, thus reducing overall losses.

LOW NITROGEN OXIDE COMBUSTION STRATEGIES

When coal is burned, oxides of nitrogen, including nitrogen dioxide, nitrogen oxide, and nitrous oxide, can be generated. Together these are generally termed *nitrogen oxide*, or NO_x. Nitrogen oxide is the product of the oxidation of the nitrogen in air with oxygen, also in the air, at the high temperatures found during coal combustion. The amount produced depends on two factors: the temperature at which the combustion takes place and the amount of oxygen available. If the amount of oxygen is restricted, then it will preferentially react with carbon rather than nitrogen, thereby reducing the production of nitrogen oxide. The rate of the reaction is also affected by the temperature; the lower the temperature the slower the reaction and the less reaction product. Both of these conditions can be exploited in a coal-fired power plant to maintain as low a level as possible in the flue gases.

One of the key strategies is to restrict the amount of oxygen available for combustion in the hottest part of the furnace (Figure 3.10). In a PC power plant, powered coal and air are generally admitted into the combustion chamber

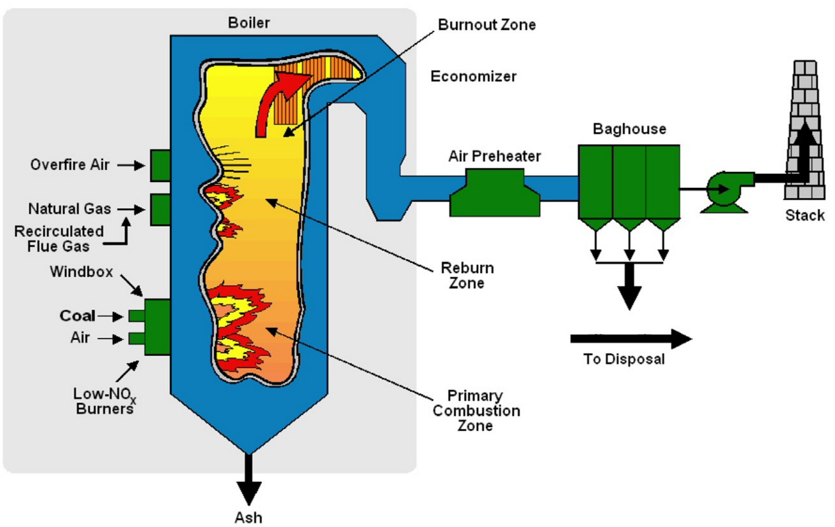

FIGURE 3.10 Low nitrogen oxide pulverized-coal boiler with reburning.

together in a continuous flow. To keep nitrogen oxide production low, some of the air needed to burn the coal completely is prevented from entering the initial combustion region with the coal; instead it is delayed briefly, being admitted to this primary combustion region after some of the combustion has been completed. This staged combustion procedure (as it is commonly known) can reduce the level of nitrogen oxide produced by 30–55%.

The initial combustion zone is normally the hottest region in the furnace. As the combustion gases leave this zone they start to cool. At this stage further air can be admitted (if combustion of the pulverized coal is still incomplete) to allow the combustion of the fuel to be completed, but at a lower temperature where the production of nitrogen oxide is reduced. The air admitted at this stage in the furnace is called overfire-air. When used in conjunction with controlled air admission in the main combustion zone, the use of overfire air can lead to a reduction in nitrogen oxide levels of 40–60%.

A third strategy that can reduce the nitrogen oxide levels even further is called reburning. This simply means that more coal, or natural gas, is introduced into the combustion gases after they have left the combustion zone. By this time gaseous oxygen levels in the flue gases will have been greatly reduced so that the new fuel will effectively steal oxygen from nitrogen oxide to combust. The effect is to remove some of the oxides of nitrogen that have been formed. Overall reductions of up to 70% can be achieved when all three strategies are applied. Low nitrogen oxide burners, overfire air, and reburning are all strategies that can be introduced into existing coal-fired power plants that do not have them, as well as being incorporated into new plants.

SULFUR DIOXIDE REMOVAL

There are no combustion strategies that can be used to control the generation of sulfur dioxide in PC plants. If sulfur is present in coal it will be converted into sulfur dioxide during combustion. The only recourse is to capture the sulfur, either before the coal is burned using a coal-cleaning process, or after combustion using some chemical reagent inside the power plant.

There are many chemicals that are potentially capable of capturing sulfur dioxide from the flue gases of a power station but the cheapest to use are lime and limestone. Both are calcium compounds that will react with sulfur dioxide to produce calcium sulfate. If the latter can be made in a pure-enough form it can be sold into the building industry for use in wall boards.

One of the simplest methods of capturing sulfur dioxide is to inject one of these sorbent materials into the flue-gas stream as it exists the furnace (Figure 3.11). Reaction then takes place in the hot gas stream and the resultant particles of calcium sulfate, and of excess sorbent, are captured in a filter downstream of the injection point.

Depending on the point of injection of the sorbent, this method of sulfur removal can capture between 30% and 90% of the sulfur in the flue-gas stream.

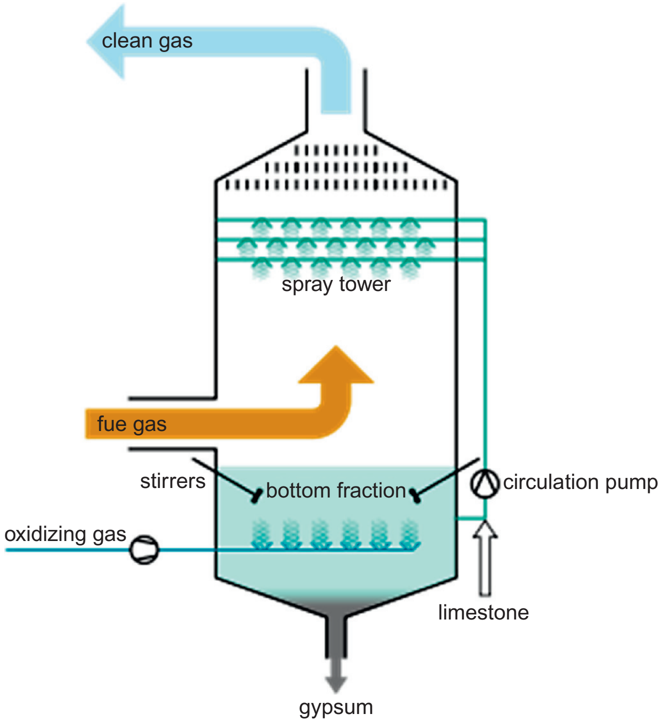

FIGURE 3.11 Schematic of a flue-gas desulfurization tower.

The cheapest, and least effective, method (30–60% capture efficiency) is to inject the sorbent directly above the furnace. Injection later in the flue-gas stream is more expensive but can remove up to 90% of the sulfur.

Sorbent injection into a flue-gas stream is one of the cheapest methods of capturing sulfur but it is not the most efficient. The best established method of removing most of the sulfur from the flue gas of a power plant is with a flue-gas desulfurization (FGD) unit, also called a wet scrubber. The FGD unit comprises a specially constructed, tall chamber through which the flue gas passes. In a typical scrubber, a slurry of water containing 10% lime or limestone is sprayed into the path of the flue gas where it reacts, capturing the sulfur dioxide. The slurry containing both gypsum and unreacted lime or limestone is then collected at the bottom of the chamber and recycled.

Typical wet scrubbing systems can capture up to 97% of the sulfur within the flue gas. With special additives, this can be raised to 99% in some cases. Wet scrubbers can easily be fitted to existing power plants, provided the space is available. Wet scrubbing technology is technically complex. It has been likened to a chemical plant operating within a power station. For this reason it requires skilled staff to operate. Nevertheless, it provides a proven method for removing high levels of the sulfur from a coal-fired power plant's flue-gas stream.

A variation on the lime or limestone scrubbing system is the seawater FGD system. Instead of the calcium-based absorbents, this uses seawater pumped directly from the sea in the scrubbing tower to absorb sulfur dioxide. Seawater is naturally alkaline and will react with the sulfur dioxide, generating soluble sulfates that are carried away with the seawater and eventually released again into the sea. Seawater scrubbing can only be used when power plants are built on coastal sites but is claimed to be capable of 98–99% removal efficiency.

Other, more complex treatment systems have been tested including the use of activated carbon or charcoal. This is capable of absorbing both sulfur dioxide and nitrogen oxide and can eventually be regenerated, leaving pure sulfur as one by-product. Activated carbon can also absorb trace elements such as mercury.

NITROGEN OXIDE CAPTURE

As with sulfur dioxide, it is possible to remove nitrogen oxide after it has been formed in the flue gas of a power plant. The process involves the injection of either ammonia gas or urea into the flue-gas stream. Either chemical reacts with the nitrogen oxide present, converting the oxides into nitrogen and water.

If the ammonia or urea is injected into the hot flue-gas stream where the temperature is between 870 °C and 1200 °C, the reaction will occur spontaneously. This is called selective noncatalytic reduction (SNCR). At lower temperatures, the reaction is too slow, so a special metal catalyst is necessary to stimulate the reaction process. Where a catalyst is used, the process is called selective catalytic reduction (SCR).

SNCR will remove between 35% and 60% of the nitrogen oxide from the flue-gas stream. However, if care is not taken it can lead to contamination of fly ash with ammonia and to ammonia slip—the release of excess ammonia into the atmosphere. Nevertheless, it has been utilized at power plants in several parts of the world.

More widely used than SNCR is SCR. SCR units are commonly found on gas turbine power stations but may also be fitted to a coal-fired power plant where low nitrogen oxide combustion strategies do not reduce the emissions levels to below the regulatory limits (Figure 3.12). Typical flue-gas temperatures in an SCR unit are 340–380 °C. The hot gases pass over a solid catalyst that is normally made from a vanadium-titanium-based material or from a zeolite. The system is generally capable of removing 70–90% of the nitrogen oxide emissions from a flue-gas stream. Ammonia slip, although still possible, tends to be less of a problem.

There are two major drawbacks to SCR. First, it can only be used with low sulfur coals (up to 1.5% of sulfur), and second, it is expensive. The catalyst requires changing every three to five years. In addition, sulfur in coal entering the SCR can be converted into trioxide (SO_3), which becomes highly corrosive on contact with water when it forms sulfuric acid. This has to be carefully controlled or it can be extremely damaging both in the power plant and in the atmosphere.

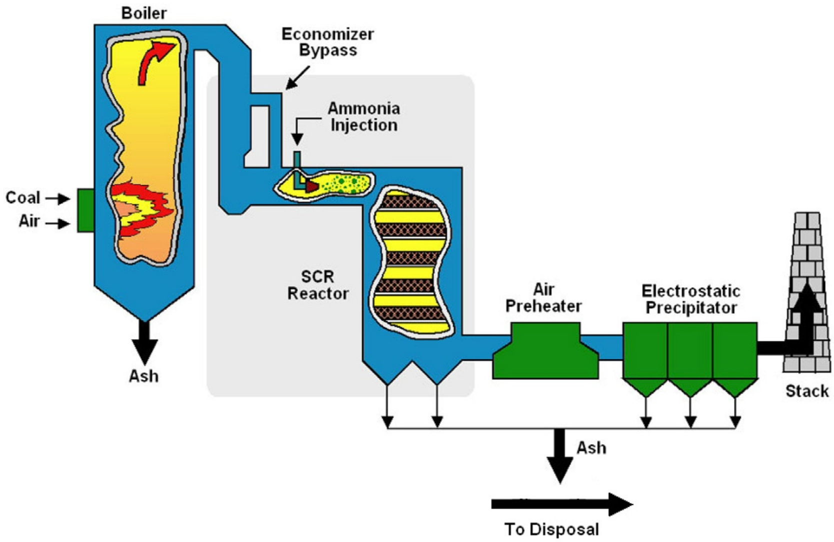

FIGURE 3.12 Coal-fired power plant with SCR reactor.

It is also possible to remove nitrogen oxide using a wet scrubbing system. The reagent sprayed into the flue-gas stream in this case is a solution sodium hydroxide, although hydrogen peroxide has also been used. The reaction that takes place leads to the formation of nitrites and nitrates that must be isolated for disposal. Wet scrubbing is not normally used in power stations for nitrogen oxide removal.

COMBINED SULFUR AND NITROGEN OXIDE REMOVAL

There are a number of processes that are capable, in principle at least, of removing both sulfur dioxide and nitrogen oxide in the same process. This ought to be more economical since only one process is needed instead of two. The use of activated charcoal has already been discussed as capable of treating both. Another is to use electron irradiation of the flue gases. Electrical discharge techniques have also been tested. However, none of these has yet been deployed widely commercially and most power plants rely on the well-tested processes described earlier.

PARTICULATE (DUST) REMOVAL

There are two principal systems for removing particulates from the flue gas of a coal-fired power station: electrostatic precipitators (ESPs) and fabric (bag house) filters.

Invented by the American scientist Frederick Cottrell, the ESP is well established and the technology has been widely exploited. It utilizes a system of plates and wires to apply a large voltage across the flue gas as it passes through the precipitator chamber. This causes an electrostatic charge to build up on the solid particles in the flue gas; as a result they are attracted to the oppositely charged plates of the ESP where they collect. Rapping the plates caused the deposits to fall to the bottom of the ESP where they are collected and removed. A new ESP will remove between 99% and 99.7% of the particulates from flue gas. However, it must be tuned to the particular coal being burned in the power plant. Where coals of different types and from various sources are to be burned, the alternative may be more effective.

Bag filters, or bag houses, are tube-shaped filter bags through which the flue gas passes on its way to the power plant stack. Particles in the gas stream are trapped in the fabric of the bags from which they are removed using one of a variety of bag-cleaning procedures. These include using supersonic blasts of air to dislodge particles so that they fall to the base of the unit and can be removed. These filters can be extremely effective, removing over 99% of particulate material. They are generally less cost effective than ESPs for collection efficiencies up to 99.5%. Above this, they are more cost effective. A system that combines a bag house–style filtration system with an ESP is under development too. This aims to provide a cost-effective high removal efficiency system, but has not yet been extensively demonstrated.

MERCURY REMOVAL

Most coals contain a small amount of mercury and this can easily end up being discharged in the flue gas from a coal-fired power plant. In the United States the emission of mercury is to be regulated and proposals have also been put forward in the EU. This will necessitate the introduction of effective capture methodologies.

Dust removal systems in power plants will remove around 25% of the mercury released during combustion. When a wet scrubbing sulfur removal system is installed too this can increase to 40–60%. Adding SCR can lead to 95% removal with bituminous coals. However, sub-bituminous coals and lignites do not respond well so alternative measures may be needed to reduce mercury emissions to below regulatory limits.

The injection of activated carbon particles has been used to remove impurities such as mercury in waste incineration plants, and this appears to offer the best solution where further mercury capture is necessary. The carbon particles will then be removed in the dust removal system through which the flue gases pass at a later stage. It is expected that plants will eventually need to remove 90% or more of the mercury released during combustion.

CARBON DIOXIDE REMOVAL

Limiting or completely removing carbon dioxide from the flue gases of a coal-fired power plant is likely to become necessary on new power plants after 2020. Even before that, increasingly onerous penalties for emitting the gas may make it necessary for power plant operators to find ways of curbing their carbon emissions (carbon emissions is often used as shorthand for carbon dioxide emissions). The technology to capture carbon dioxide already exists. It has not been deployed on a large coal-fired power plant but demonstration projects are being planned. However, this is only part of the problem. Carbon dioxide, once captured, must be stored (sequestered) in a way that prevents it ever entering the atmosphere. Otherwise the capture process will have been a waste of time. Therefore, carbon sequestering techniques must be perfected alongside carbon capture if the ultimate goal of carbon-emission-free coal combustion is ever to be realized.

For carbon capture there are three main approaches available. The first, often called post-combustion capture, involves installing a plant similar to an FGD scrubber to the exhaust of the power plant. Reagents capable of absorbing carbon dioxide are already available and this technology is likely to be one of the simplest and possibly the most economical to deploy. An alternative to this, called pre-combustion capture, involves pretreating coal to remove the carbon before combustion. This is achieved using a modified version of coal gasification that, when carried out fully, leaves a fuel gas composed primarily of hydrogen. The gasification process also requires some form of carbon dioxide capture technology, which may be via a scrubber system too. The third way of tackling carbon dioxide emissions attempts to sidestep the difficult problem of separating carbon dioxide from the flue gases, which are mainly composed of nitrogen after combustion is over. Instead, oxygen is separated from air first, and then coal is burned in virtually pure oxygen. The scheme, called oxy-fuel combustion, leads to an exhaust gas stream composed primarily of carbon dioxide, mixed with some water vapor and excess oxygen, from which it is much easier to isolate the carbon dioxide than when the latter is mixed with nitrogen.

There is a further method of reducing, though not eliminating, carbon dioxide emissions in a coal-fired plant: by replacing some of the coal with a biofuel. Biofuel cofiring, as this technique is known, cannot be used to eliminate all carbon emissions. For that to be possible the plant would have to burn biofuel alone, so it would no longer be a coal-fired plant. However, it can improve a coal plant's environmental performance.

Post-combustion Capture

Biomass as a power plant fuel is generally considered to be carbon dioxide neutral. Combustion of biomass generates carbon dioxide in the same way as the combustion of coal but provided the biomass is derived from a source that is constantly replenished by regrowth. Then the new biomass that grows will absorb carbon dioxide from the atmosphere equivalent to that released during combustion.

Coal-fired power plants will burn biomass in exactly the same way as they burn coal. Depending on the amounts involved, the plant may require modification. Small quantities of biomass, up to perhaps 7%, can be simply mixed with the coal that is used to fire the plant without any adaptation being necessary. Pulverized coal and pulverized biomass enter the combustion chamber together and there is little change to operating conditions.

For the proportion of biomass to be increased beyond this, dedicated biomass burners must be used so that the conditions within the furnace can be more carefully controlled. With this technique up to 20% biomass can be introduced into a modern supercritical PC plant without adversely affecting its operation. This is also an advantageous means of using biomass because the efficiency of a modern coal-fired power plant will be much higher than of a dedicated biomass plant.

It is also possible to gasify biomass, generating a low-calorific-value fuel that can also be burned in a furnace. Introducing this gas rather than solid biomass into a coal-fired plant may offer some operational advantages. However, it does require a gasifier, making it more expensive and technically more complex than simple biomass cofiring.

Coal Gasification

The most straightforward approach to eliminating the carbon emissions resulting from coal combustion in a power station is to capture the carbon dioxide produced in a chemical absorption plant. This is similar in concept to an FGD scrubber for removing sulfur dioxide. Carbon dioxide capture is carried out today in a number of industries using a reagent called monoethylamine (MEA) and the process has even been applied on a small scale at a number of power plants to provide pure carbon dioxide for the foods industry. Ammonia and other amine reagents can also be used to capture carbon dioxide and this is a fertile area for research with a wide range of alternative absorbents and techniques being studied. These include cryogenic capture, solid rather than liquid absorbents, and membrane systems.

The key to any reagent for carbon dioxide capture is that the chemical reaction must be easily reversible so that the carbon dioxide can be released again and the reagent regenerated. Ease of capture depends on the concentration of carbon dioxide, which in the flue gases of a typical PC power plant will be around 14% by volume. At this concentration some form of chemical absorption process is necessary to achieve the required level of capture efficiency. This makes the regeneration more energy intensive than would be the case with a physical absorbent. Amines appear to be the best reagents available today but others may become available in the future.

In a typical post-combustion capture plant the flue gases from the boiler are passed up a tall tower into which the reagent is sprayed so that gases and the liquid absorbent mix thoroughly (Figure 3.13). Depending on the column size, around 90% of the carbon dioxide within the flue gases can be removed in this way. The MEA reagent that has been sprayed into the path of the flue gases is collected at the bottom of the tower and then pumped to a second unit that heats it, releasing the carbon dioxide and returning pure MEA to enter the scrubbing tower once again. The regeneration process is extremely energy intensive, taking around 23% of the power output of the power station. A further 8% of the

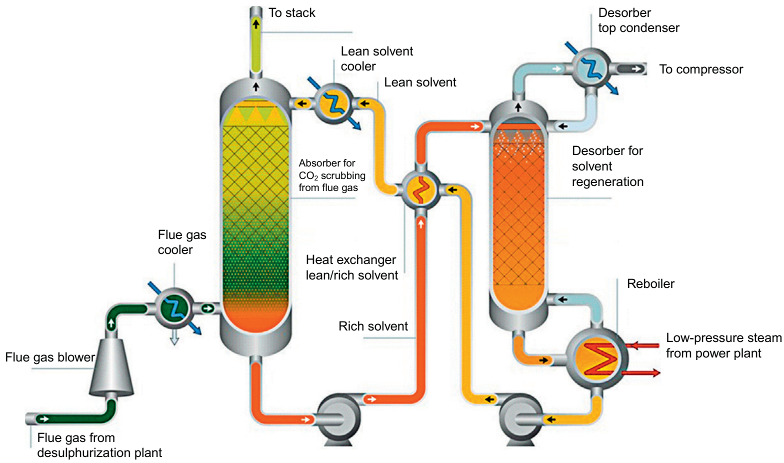

FIGURE 3.13 Post-combustion carbon dioxide capture system using absorption tower.

output is then needed to compress the carbon dioxide so that it can be transported for sequestration, and the MEA absorption plant uses a further 5% of the plant output as auxiliary power. This reduces the overall efficiency of power generation significantly. For example, an ultra-supercritical PC power plant that might achieve an efficiency of 45% would have this cut to around 35% with post-combustion capture. This has a large effect on the cost of electricity from such a plant.

Aside from its relative simplicity, one of the attractions of post-combustion carbon capture is that it can be applied retrospectively to existing power stations. This is unlikely to be cost effective for a large part of the old fleet of subcritical coal-fired power plants operating today, but it could be economical to retrofit newer and more efficient supercritical plants.

Oxy-fuel Combustion

When coal is burned in air, the carbon dioxide resulting from the combustion reaction is mixed with a large quantity of nitrogen along with a number of other constituents of air, pollutant gases, and particles generated during combustion. The carbon dioxide accounts for only around 14% of the mixture and removing this low concentration is relatively difficult. If, on the other hand, coal is burned in pure oxygen, then the flue gases resulting will be almost pure carbon dioxide. This will be mixed with the little nitrogen left over after oxygen purification, some water vapor, and the various pollutants generated. Isolating carbon dioxide from this mixture is much simpler than in the case of air combustion. This is the concept behind oxy-fuel combustion.

Oxy-fuel combustion essentially swaps a carbon dioxide separation plant for an oxygen separation plant so the efficiency of this plant now becomes key to

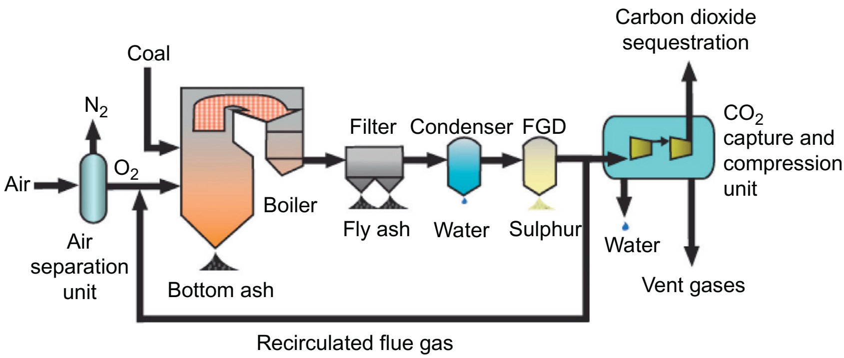

FIGURE 3.14 Schematic of an oxy-fuel combustion power plant.

the overall efficiency (Figure 3.14). Oxygen separation plants are widely used industrially but it is an energy-intensive process. Reducing the energy cost is likely to be important if the technology is to become successful.

Oxy-fuel combustion is not simply a matter of replacing the air with oxygen for combustion in a coal-fired plant. Other changes must be made too. When coal burns in oxygen the flame temperature can reach 2500 °C, far higher than current materials can stand. To overcome this, some of the carbon dioxide–rich flue gas from the exhaust of the boiler must be fed back and mixed with the oxygen, diluting it and reducing the flame temperature in the process.

Energy conversion efficiency for an oxy-fuel combustion PC plant appears to be slightly lower than for a plant with post-combustion capture; estimated efficiencies are in the range 29–32%, but since the technique has not yet been tested at a commercial scale, this remains speculative. However, like post-combustion capture, it is feasible to retrofit oxy-fuel combustion to an existing coal-fired power plant. Again this is only likely to be effective with a high-efficiency supercritical plant, but it does offer an interesting alternative for future carbon capture.

Biomass Cofiring

As noted before, coal gasification is another well-established industrial technique, used both as a source of domestic gas in the past and in the chemical industry to convert coal into a gas called synthesis gas or syngas that can be used as a precursor for a variety of chemical processes (Figure 3.15).

Coal gasification is carried out in a reactor called a gasifier where the fuel is burned under carefully controlled conditions with a strictly limited amount of either air or oxygen, together with steam. Oxygen gasifiers require a dedicated oxygen plant. The partial combustion of the coal, according to the reaction shown earlier, results in a gas that is primarily a mixture of hydrogen and carbon monoxide together with various combustion impurities such as hydrogen sulfide. A residual char is left.

Syngas still contains carbon in the form of carbon monoxide, so to create a carbon-free gas this must be passed through a second reactor where it is mixed with further steam and passed over a catalyst that converts the carbon monoxide into carbon dioxide and further hydrogen:

$$CO + H_2O = CO_2 + H_2$$

This process is called the water-shift reaction.

This mixture of carbon dioxide and hydrogen must now be cleaned and then the carbon dioxide separated. The latter is in much higher concentration than in the flue gases of a coal-fired power plant, accounting for around 50% of the mixture, and it can be separated using a physical absorbent rather than one that reacts chemically with the gas. The physical absorbent binds the carbon dioxide more loosely than a chemical absorbent such as MEA would, so releasing the carbon dioxide and regenerating the solvent is much less expensive from an energetic point of view. Reagents based on glycol and refrigerated methanol are used industrially to separate hydrogen from carbon dioxide, and these could be applied in a power plants based on gasification too.

To produce electricity from coal in this way, the hydrogen is burned in a gas turbine–based combined cycle power plant that is closely integrated with the gasification process to achieve the highest efficiency possible. This configuration is a variation of the IGCC power plant already described. As with the other approaches to carbon dioxide capture and removal, this leads to a significant fall in efficiency, and an IGCC plant with carbon capture appears likely to have an

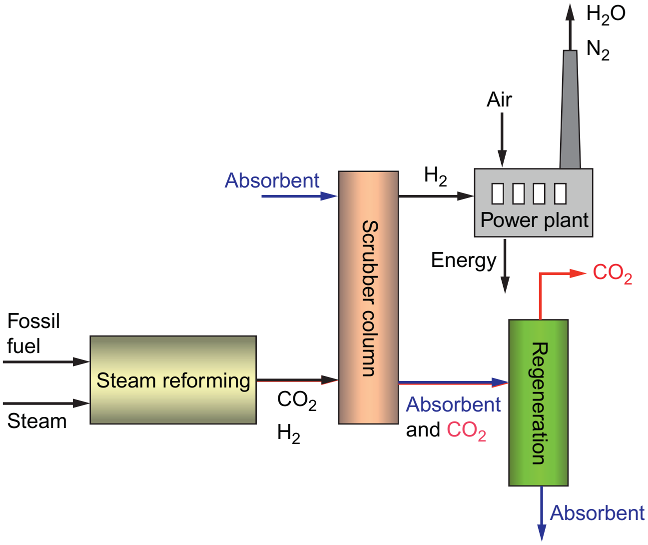

FIGURE 3.15 Schematic of a power plant with pre-combustion capture for carbon removal.

energy conversion efficiency of around 32%, broadly similar to the other technologies discussed earlier.

An alternative to an IGCC plant is to use the hydrogen produced by the gasification process as fuel for a fuel cell. These devices, which are potentially capable of high efficiency when operating with pure hydrogen, will be discussed in a later chapter.

CARBON DIOXIDE SEQUESTRATION

If carbon dioxide is to be captured during the use of coal (or any other fossil or carbon-based fuel) for electricity generation, then it must also be stored in such a way that it cannot ever enter the atmosphere. Otherwise the capture is pointless.

Various methods are being explored for achieving this but the most likely solution is that the carbon dioxide will be stored underground (Figure 3.16). Several options exist for this. One of the most obvious is to pump it into exhausted oil and gas fields. The infrastructure already exists to carry this out and carbon dioxide has been pumped into older oil fields to enhance oil recovery for many years so the technology is well understood. Questions remain about the security of the gas once it has been sequestered in this way, but in principle it could provide a safe means of storing the gas.

Exhausted wells offer a first choice for testing capture and sequestration technology, but they cannot provide the capacity that will be needed if this is to become a global practice. For that, underground geological formations such as aquifers are likely to be used. These are sites where an impermeable rock cap

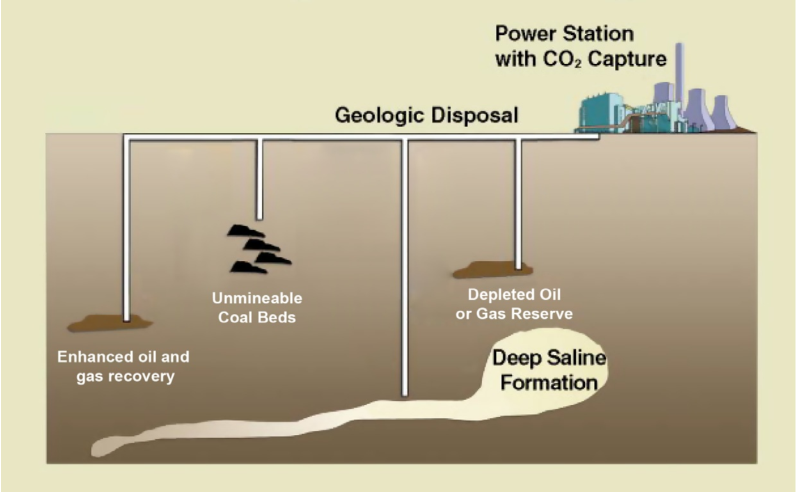

FIGURE 3.16 Carbon sequestration.

traps liquids beneath it. Carbon dioxide can be pumped into such aquifers where it first displaces the brine already trapped there and eventually dissolves in it, potentially forming solid salts that are permanently held.

Some consideration has also been given to storing carbon dioxide in the world's oceans. These are already the largest global carbon dioxide sinks, holding around 20 times as much carbon dioxide as there is in the atmosphere and biosphere combined. However, there are too many uncertainties attached to this method of carbon dioxide disposal for it to be used today.

Once carbon sequestration sites have been found and stores established, then the carbon dioxide must be delivered to them. In most cases this is likely to be via a new carbon dioxide pipeline system. The gas must first be pressurized and then it can be pumped from power plants to the storage sites. The gas is normally pressurized between 10 MPa and 15 MPa, when it enters a supercritical state under which it has the density of a fluid though it can still behave like a gas. Tanker transportation might be feasible under some circumstances but it does not appear as economical as pipeline transportation.

COST OF COAL-FIRED POWER GENERATION

Coal-fired power stations are relatively expensive to build since their construction involves both large quantities of expensive materials, such as iron and steel, and large volumes of labor. While some parts of a coal-fired power plant such as its steam turbines can be assembled in a factory and then delivered to the site, much of the assembly of the boiler and flue-gas cleaning systems must take place at the site itself. As a consequence, the cost of a coal-fired power plant will be vulnerable to changing commodity costs and generally increasing labor costs.

Against this the fuel (coal) is generally the cheapest of fossil fuels and this will normally outweigh the high capital cost so that the cost of electricity from a coal-fired power station will be among the most competitive available. Currently, coal-fired power stations are built without carbon capture facilities but this is likely to be required at some point during the third decade of the 21st century. It is also likely that plants built before that time will be required to be carbon-capture ready so that post-carbon-capture technology can be fitted at a later stage. All this needs to be taken into account when considering construction of a coal-fired power generating facility.

Table 3.6 shows the estimated 2011 cost of coal-fired power plants in the United States in 2010 dollar prices based on data from the U.S. government's Energy Information Administration. This analysis suggests that the capital cost of an advanced PC power plant is \$3167/kW. When carbon capture and storage is added to this, based on post-combustion capture, the cost rises to \$5099/kW, an increase of 60%.

An integrated gasification combined cycle power plant without carbon capture and storage would, on the same basis, have cost \$3565/kW in 2011.

TABLE 3.6 Cost of Coal-fired Power Plant Technology

Type of Coal-fired Power Plant	2011 Capital Cost ($/kW)
Advanced pulverized coal plant	3167
Advanced pulverized plant with carbon capture and storage	5099
Integrated gasification combined cycle plant	3565
Integrated gasification combined cycle plant with carbon capture and storage	5348

Source: *Updated Capital Cost Estimates for Electricity Generation Plants,* U.S. Energy Information Administration, 2010.

TABLE 3.7 Capital Costs in 2015 and 2025 for Various U.S.-generating Technologies (2010 U.S.$)

Type of Coal-fired Power Plant	2015 ($/kW) without Carbon Dioxide Capture	2025 ($/kW) Carbon Capture and Storage
Pulverized coal-fired plant	2000–2300	2600–2850
Integrated gasification combined cycle plant	3150–3450	3100–3800

Source: Program on Technology Innovation: Integrated Generation Technology Options, Technical Update June 2011, Electric Power Research Institute, 2011. (EPRI reference 1022782.)

However, when carbon capture and storage is added to this plant, based on the type of system discussed before, the price would rise to $5348/kW. The cost of an oxy-fuel combustion plant was not included in the analysis but other analyzes suggest that the cost is likely to be similar to those in Table 3.6 for the other technologies with carbon capture.

Other analyses suggest that future costs will be lower than this. Table 3.7 shows another set of figures for the cost of future coal-fired power plants, in this case as estimated by the U.S. Electric Power Research Institute (EPRI). This analysis is based on the assumption that carbon dioxide capture will not be deployed by 2015. The cost of a PC power plant in that year, again calculated in 2010 U.S. dollars, is estimated to be between $2000/kW and $2300/kW. At the same time, the cost of an integrated gasification combined cycle plant is put at between $3150/kW and $3450/kW.

EPRI has assumed that carbon capture and storage will be necessary in 2025. By then the cost of a PC plant with carbon capture and storage is estimated to be

\$2600–2850/kW. For an integrated gasification combined cycle power plant the cost range is \$3100–3800/kW. As before, there is no consideration of an oxy-fuel combustion plant, but the cost of this can be expected to fall somewhere within the cost range of the two other technologies, \$2600–3800/kW.

All these costs are for power plants built in the United States. While costs are likely to be similar in Europe, commodity prices and especially labor costs are likely to differ widely in other parts of the world and this can affect the capital cost significantly. Costs in countries like China and India are consistently much lower than in the United States and Europe.

Chapter 4

Natural Gas–fired Gas Turbines and Combined Cycle Power Plants

Power stations that use gas turbines as their primary energy conversion system have in recent years become one of the mainstays of the power generation industry in the developed world and have also made inroads in the developing world. These high-technology engines can burn a range of liquid and gaseous fuels but most commonly they are fired with natural gas and their success depends primarily on its cost and availability.

The current prominence of gas turbine technology is a striking position to have achieved for a technology that until the late 1980s played a small part in global power generation. Several factors have conspired to bring about its success. First was the recognition of natural gas as a valuable fuel rather than a by-product of oil production, which was best disposed of by flaring at the well head. Though some natural gas is still flared, the fuel is now a vital part of the global energy economy, particularly through its use for power generation.

The cost of natural gas was a second factor in its initial rise in popularity. This was particularly true in the early 1990s when it was cheap in many developed markets. Since then the price of natural gas has become one of the most volatile among power plant fuels and the economic viability of gas turbine power generation has waxed and waned with its cost.

A third, and perhaps the overriding, factor in the success of gas turbine technology has been the development and continual rise of the gas turbine combined cycle plant. Combined cycle plants are one of the cheapest types of power-generating facility to install and the best modern examples can demonstrate energy conversion efficiencies of 60%, higher than any other large-scale fossil fuel–fired power plant in operation today. This, together with the fact that the combustion of natural gas produces significantly less carbon dioxide and less of most other pollutants than coal combustion, has made the gas turbine an attractive addition to the armory of electricity-generating technologies.

The role of gas turbines may become more important still. The International Energy Agency (IEA) suggested in 2011 that the world could be entering a golden age of gas, a development that would benefit gas-fired power generation significantly. The IEA proposition depends on the coincidence of the pressing global need to reduce carbon dioxide emissions,[1] the concomitant rise in the use of fluctuating sources of renewable generation such as wind and sun, and the sudden improvement in natural gas availability brought about by shale gas exploitation.

The ability to exploit shale gas economically has led to a sharp rise in the availability of natural gas in the United States and a consequent fall in the cost of the fuel. The trend is expected to spread to other regions as shale gas exploration expands. At the same time, the flexibility of combined cycle power plants burning natural gas makes them an ideal alternative source of power to help balance the varying generation from renewable sources. Thus, the use of wind power and solar power, together with gas-fired combined cycle power plants, is seen as a compelling combination for the future of power generation in many regions of the world.

Taken together all these factors could see the use of natural gas as a source of electricity double between 2008 and 2035. Based on figures from the U.S. Energy Information Administration it could account for 24% of global electricity generation in 2035. While this would be lower than the proportion of generation from coal, it would represent slightly more than the total from all renewable sources.

NATURAL GAS

While gas turbines are capable of burning a range of fuels, including distillate fuel oil, hydrogen, and gases produced by gasification of both coal and biomass, the main fuel for which the majority of gas turbine power plants are built is natural gas. Natural gas as it is extracted from gas fields is a mixture of combustible hydrocarbons. The main component is methane (CH_4), which normally accounts for 70–90% of the total.[2] Other hydrocarbons such as ethane, propane, and butane can account for up to 20% of the mixture and there may be up to 8% carbon dioxide, small amounts of oxygen and nitrogen, and up to 5% hydrogen sulfide. The gas is normally cleaned after it has been pumped from the ground to remove impurities such as hydrogen sulfide (this can be processed into pure sulfur), carbon dioxide, and water. The higher hydrocarbons, such as propane and butane, may also be removed for industrial use and the cleaned gas, now referred to as dry natural gas,[3] is ready for use.

Natural gas is found in many parts of the world but most countries have only small reserves. Regionally, the Middle East has the largest total proven recoverable reserves at 80 trillion m^3 or 38.4% of the global total in 2011, as shown in Table 4.1. The other major reserves are in Europe and Eurasia with

1. Switching from coal to natural gas combustion reduces carbon dioxide emission rates significantly.
2. The gas used for household heating and cooking is almost pure methane.
3. Straight from the well it is referred to as wet natural gas.

TABLE 4.1 Proved Recoverable Natural Gas Reserves

	Reserves (trillion m^3)	Percentage of Total	Reserve/ Production Ratio
North America	10.4	5.2%	12.5
Central and South America	7.6	3.6%	45.2
Europe and Eurasia*	78.7	37.8%	75.9
Middle East	80	38.4%	—
Africa	14.5	7%	71.7
Asia Pacific	16.8	8%	35
Total	187.1	100%	63.6

**The figure includes reserves of 44.6 tn m^3 in Russia.*

Source: BP Statistical Review of World Energy.

78.7 trillion m^3 or 37.8% of the global total. Note, however, that these reserves are dominated by Russian gas, which accounts for 21.4% of the global total. Aside from these two regions, the others in Table 4.1 each hold less than 10% of global reserves.

When broken down by country, four nations hold the bulk of global reserves. The Russian Federation on its own has 44.6 trillion m^3 of proven reserves, 21.4% of the total, and Turkmenistan has 24.3 trillion m^3, 11.7% of the total. In the Middle East, Iran holds a further 33.1 trillion m^3 (15.9%) and Qatar 25 trillion m^3 (12%). Among them these four command 61% of proven global natural gas reserves. Other nations rich in natural gas include the United States with 8.5 trillion m^3, Venezuela with 5.5 trillion m^3, Saudi Arabia with 8.2 trillion m^3, and the United Arab Emirates, Kuwait, Algeria, Nigeria, and Australia.

Natural gas consumption has been rising steadily over the past four decades (Table 4.2). From 651 billion m^3 in 1965, consumption had reached 1649 billion m^3 in 1985 and 2782 billion m^3 in 2005. By the end of 2011 global annual consumption was 3223 billion m^3. The largest national consumer, by a wide margin, is the United States, which used 690 billion m^3 in 2011, followed by the Russian Federation with 427 billion m^3.[4] China, by comparison, consumed only 131 billion m^3.

Natural gas is readily transported by pipeline, so this has become the preferred method where pipeline construction is possible. The United States has

4. The cost of natural gas is heavily subsidized in Russia, as it is in Iran, another heavy user.

TABLE 4.2 Global Natural Gas Consumption

Year	Consumption (billion m^3)
1965	651
1970	987
1975	1186
1980	1437
1985	1647
1990	1960
1995	2135
2000	2412
2005	2782
2010	3169
2011	3223

Source: BP Statistical Review of World Energy.

extensive pipeline networks, as does Europe, while pipelines from central Asian countries and Middle East nations rich in the resource are beginning to extend both east and west.

As a consequence, much natural gas is traded internationally through pipelines. However, where pipeline delivery is not possible, transportation in liquefied form is possible. Major gas producers such as Qatar liquefy gas for delivery by ocean-going tankers to many parts of the world. Important users of liquefied natural gas (LNG) include Japan, Taiwan, and, more recently, the United States and United Kingdom. With its new shale gas resources, the United States is expecting to become an exporter of LNG in the near future.

Global natural gas supplies are sufficient for another 63.6 years, overall, as the figure in the final column of Table 4.1 indicates. However, local national supplies in some regions are under much greater stress. Most striking is North America where existing proven reserves will last only another 12.5 years at current rates of consumption.[5]

5. This depends on how the reserves are estimated. If potential reserves are taken into account the lifetime is much greater. For example, other estimates have suggested current U.S. reserves will last for between 80 years and 120 years.

GROWTH OF GAS TURBINE TECHNOLOGY

The gas turbine was originally developed during the 1930s and 1940s as an aviation engine and it remained almost exclusively the preserve of the aviation industry until the end of the 1960s. The potential for gas turbines to be used in power generation began to be exploited during the 1970s and early 1980s when they entered service for standby and peak power support on national grids. These early aero-derivative gas turbines were broadly similar in design to the aero engines upon which they were based and shared many components. They were light, they could start up quickly, and they were able to change power output rapidly, making them ideal for a grid support role.

Recognizing the value of gas turbines to the power generation industry, some manufacturers began to design industrial gas turbines intended specifically for this market. These were heavier in construction than their aero-engine counterparts since weight was not an issue for stationary applications such as power generation. As their design evolved, and as their performance was tailored specifically for the power generation market, these industrial gas turbines moved further away from aero engines. At the same time, aero-derivative stationary engines remained in production alongside the new industrial ranges of products.

It was toward the end of the 1980s that the first big combined cycle power plants were built using gas turbines. These used one or more gas turbine generators to produce electricity, with the heat from the turbine exhaust utilized to raise steam in a special heat-recovery boiler. The steam was then used to generate more electrical power in a steam turbine generator. This combination led to a high-performance power plant configuration that, by 1990, was capable of around 50% energy conversion efficiency.

The market for combined cycle power plants grew rapidly during the 1990s, particularly in Europe and the United States where companies operating in the newly liberalized and deregulated electricity markets found the low cost of such plants appealing. Steep rises in the cost of natural gas toward the end of the decade blunted the appeal somewhat. Even so they continued to be popular during the first decade of the 21st century, although gas price volatility during the decade made their economics sometimes questionable. In the second decade of the century with gas availability predicted to rise and prices expected to stabilize, the economics of combined cycle generation look promising again.

To cater for the growing power generation market, the main manufacturers—and there are only a limited number of these because the gas turbine is a very specialized high-technology machine—quickly began to build bigger and more efficient engines. Today, single engines of up to 400 MW in capacity are available. Most of these are aimed at the combined cycle market and have efficiencies on their own between 38% and 42%. However, when used in a modern integrated combined cycle plant the latter should be capable of around 60% efficiency.

GAS TURBINE PRINCIPLE

A gas turbine is a machine that harnesses the energy contained within a gas—either the kinetic energy of motion of a flowing gas stream or the potential energy of a gas under pressure—to generate rotary motion. In the case of a gas turbine this gas is usually, though not necessarily, air. The earliest human-made device of this type for harnessing the energy in moving air was a windmill described by Hero of Alexandria in the 1st century AD.

This earliest known windmill was a near relative of today's wind turbine, which though clearly a type of gas turbine, is far removed from the modern gas turbine concept. Closer in concept to the gas turbine was the smokejack, developed in the middle of the 2nd century AD. As described in the 17th century by John Wilkins, Bishop of Chester, the smokejack used hot air rising through a chimney to move windmill vanes within that chimney and drive a shaft that could be used to rotate a spit for roasting meat.

This principle of harnessing moving air within an enclosed chamber to create rotary motion for driving machinery was developed further during the industrial revolution. Following this principle, the 19th century saw a number of predecessors to the gas turbine. These used some form of compressor to generate a human-made flow of pressurized air that was fed into an enclosed turbine. In these machines the compressor was usually separate from the turbine.

The direct ancestor of the modern gas turbine was first outlined in a patent granted to German engineer F. Stolze in 1872. In Stolze's design, as in that of all modern gas turbines, an axial compressor was used to generate a flow of pressurized air. This air was then mixed with fuel and ignited, creating a very energetic flow of hot, high-pressure gas that was fed into a turbine. Crucially, the compressor and the turbine were mounted on the same shaft (Figure 4.1).

Whereas a gas turbine supplied with pressurized gas from a separate compressor must inevitably rotate provided it has been designed correctly, the arrangement patented by Stolze need not necessarily do so. This is because the energy to operate the compressor that provides the pressurized air to drive the turbine is produced by the turbine itself. Thus, unless the turbine can generate more power than is required to turn the compressor—the energy for this being provided by the combustion of fuel in the compressed air that same compressor has

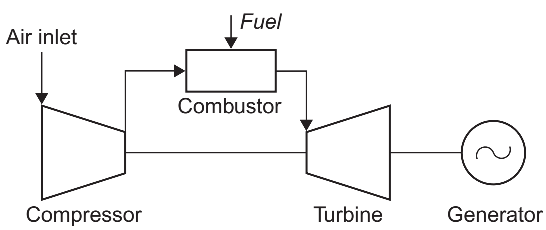

FIGURE 4.1 Block diagram of a gas turbine for power generation.

produced—the machine will not function. This, in turn, demands efficient compressors and turbines. Both need to operate at a minimum efficiency of around 80%. Only if this condition is met will the turbine operate in a continuous fashion.

The turbine system described by Stolze, although envisaging virtually all the features of a modern gas turbine, was not capable of sustained operation because the machinery sophisticated enough to achieve it had not yet been developed. The first machine that was able to operate in a sustained fashion was built in Paris in 1903. This, though, did not have a rotary compressor on the same axis as the turbine. That honor fell to a machine built by Aegidus Elling in Norway and operated later in 1903. In Elling's machine the inlet gas temperature was 400 °C, high for its time though much lower than gas turbine inlet temperatures today.

Development of the gas turbine based on this principle continued through the early years of the 20th century with the aim of generating either a flow of compressed air, rotary motion, or both for industrial use. Then, during the 1930s, the potential of the gas turbine to provide the motive force for flight was recognized and aircraft with jet engines based on the gas turbine were developed in Germany, Great Britain, and the United States. These led, in turn, to the modern aircraft engines that power the world's airline fleets.

During the late 1970s and early 1980s, as already noted, gas turbines began to find a limited application in power generation, and during the 1990s this use expanded with development of combined cycle power plants.

MODERN GAS TURBINE DESIGN FOR POWER GENERATION

A modern gas turbine maintains the three principle components of the Stolze design: a compressor, a combustion chamber, and a turbine stage (Figure 4.2). As in that design, these are closely coupled with the two rotary components on the same axis. The main energy-producing component is the turbine stage

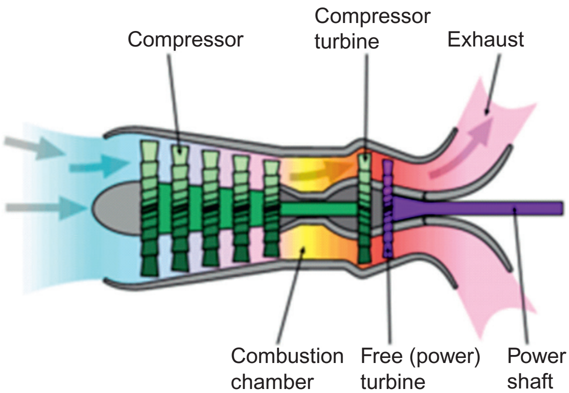

FIGURE 4.2 Cross-section of a gas turbine.

that, as outlined earlier, drives its own compressor as well as provides energy to turn a generator and produce electricity.

The turbine stage of a gas turbine is a thermodynamic heat engine. As with all heat engines, the efficiency of its operation depends on the amount of energy it can extract from the operating fluid (in this case air). For an ideal heat engine this in turn depends on the temperature and pressure difference between the gas entering the engine inlet and the gas at the turbine exhaust.

It is in the operation of the turbine stage that aero engines and power generation turbines diverge. An aero engine only needs the turbine to produce sufficient rotary power to drive its compressor. All the rest of the energy within the gas should be delivered through the turbine exhaust as a high-pressure flow of gas providing thrust. For a power generation turbine, as much energy as possible should be captured in the turbine rotary motion and as little energy as possible should exit through the turbine exhaust, which should release gases at the lowest temperature and pressure possible.

The input energy for the turbine depends on both temperature and pressure of the air at the turbine inlet. The pressure is a design feature that will vary from manufacturer to manufacturer and is defined by the compression ratio chosen for the design. All modern gas turbines utilize axial compressors that draw in and compress air. These comprise several stages of blades (much like a series of windmills but operating in reverse) to compress air to between 15 and 30 times atmospheric pressure, depending on the design (i.e., the compression ratio is between 15:1 and 30:1). A modern unit might have 10–12 sets of compressor blades (also called stages). Efficiency is typically 87%.

High-pressure air from the compressor then enters a combustion chamber where it is mixed with fuel and ignited, increasing the temperature of the air to as high as 1600 °C. At temperatures this elevated, nitrogen from air is easily oxidized to produce nitrogen oxides, which will then appear in the gas turbine exhaust. Since nitrogen oxide emissions must be controlled, gas turbine combustion chambers are designed to maintain reducing conditions as far as possible, with staged combustion to limit the production of nitrogen oxide. In some cases water is injected into the combustion chamber to reduce nitrogen oxide levels but most large turbines use dry combustion chambers.

Combustion chambers come in a variety of designs and dispositions. In some gas turbines they are kept separate from the turbine body. Other designs position them within the body, between compressor and turbine stages, while in others there are multiple combustion chambers arranged annularly around the body of the turbine. At least one manufacturer uses sequential combustion similar in concept to reheat in a steam turbine. In this case the turbine stage of the unit is split into two parts and the hot gases exiting the first part pass into an additional combustion chamber where further fuel is ignited to increase temperature and pressure again before entering the second part.

The turbine stage of a modern gas turbine will normally comprise three to five stages of blades (windmills operating as windmills in this case) operating with an efficiency of around 89%. Some designs have both compressor and

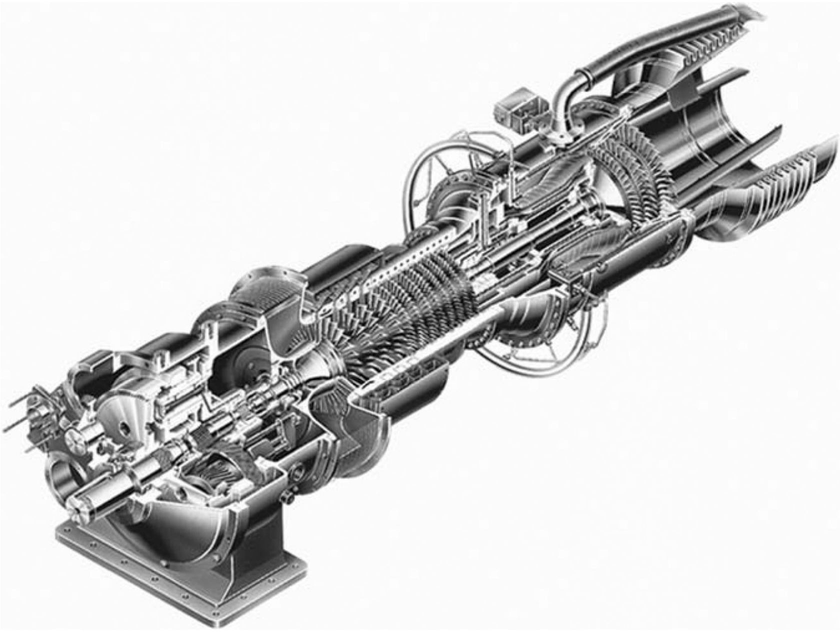

FIGURE 4.3 Cross section (photograph) of a gas turbine. *Source: Courtesy of Solar Turbines Incorporated.*

turbine blades mounted rigidly onto the same shaft (Figure 4.3). In others there are two concentric shafts—one carrying the compressor blades and the first one or two turbine stages. These turbine stages power the compressor while the latter stages, on a second shaft, are attached to a generator and produce power. Some aero-derivative gas turbines take this complexity a stage further with the compressor also divided into two parts too—a low-pressure compressor and a high-pressure compressor. The low-pressure compressor is driven by the low-pressure turbine stage of the unit and, via a concentric shaft, the high-pressure turbine drives the high-pressure compressor.

The exhaust gases from a simple gas turbine generator are released to the atmosphere. For the most efficient aero-derivative gas turbine, the outlet gas temperature will be around 400 °C to 500 °C. This represents a significant energy loss. Even so, such units can be relatively efficient. In fact, these aero-derivative gas turbines are usually the most efficient of all gas turbines with the best recording an energy conversion efficiency of 46%. Power output from this type of turbine is usually under 100 MW. Small industrial turbines tend to have slightly lower efficiencies, up to 42%.

For an open-cycle gas turbine (i.e., one not operating in a combined cycle plant), lower exhaust gas temperature equates to higher efficiency, which in turn equates to better economy. In a combined cycle plant, on the other hand, a higher exhaust gas temperature is usually preferable because this allows the second part of the plant, based on a steam turbine, to operate more efficiently. Thus, most of the largest industrial gas turbines that are principally designed for combined cycle operation will actually have a lower efficiency in open-cycle mode than some smaller turbines to obtain the best efficiency in combined cycle mode. The largest industrial gas turbines have efficiencies between 38% and 42%.

GAS TURBINE DEVELOPMENT

Efficiency is perhaps the most important operating parameter for all gas turbine operations. In aero engines, higher efficiency equates to lower fuel costs. Similarly in the power industry, higher efficiency leads to a lower unit cost for electricity generation. Efficiency is important from an environmental perspective too, because the more efficient a gas turbine–based power plant is at generating electricity, the less carbon dioxide and other atmospheric pollutants the power plant will create for each unit of electricity it produces.

As a consequence, the main focus of gas turbine development over the past 20 years had been primarily aimed at increasing efficiency. Overall, efficiency will depend on the efficiency of both the compressor and the turbine, but with both of these now highly efficient, the main means available to gas turbine designers to improve efficiency are by increasing temperature and pressure of the gas entering the turbine (i.e., heat engine) stage of the unit.

The engine compression ratio (the amount by which the compressor increases the inlet air pressure) is one variable available to designers to modify as they seek higher efficiency, but the chosen ratio is generally an optimum for a particular design rather than the highest possible. As already noted, compressor design has reached a high level of sophistication and modern compressors can deliver whatever pressure is required within the stationary gas turbine range of 15:1 to 30:1 at high efficiency.[6]

The design of the turbine stage has also reached a very high level of technical sophistication, with the shapes of the rotating blades and the stationary vanes (sometimes called nozzles) optimized using computer modeling. Thus, the only major variable still open to improve overall efficiency is turbine inlet gas temperature.

Gas turbine combustors are capable of delivering gas at higher temperatures that are currently in use, but the maximum temperature that can be exploited is limited by the performance of the materials employed in the early stages of the gas turbine. Inlet temperatures at the first stages of gas turbines have risen steadily, from around 900 °C in 1967 to 1425 °C in 2000 and 1600 °C in 2010. Efforts to raise temperatures further are underway and a Japanese program is aiming to achieve 1700 °C in the near future.

Such high temperatures place extreme demands on the materials used in turbine construction, which have to withstand these extreme conditions. The turbine blades, the vanes that control air flow from one set of blades to the next, and

6. Modern aero engines may have compression ratios of up to 40:1.

other hot-gas-path components are commonly made from nickel-based alloys cast in single crystal form to increase their resistance to deformation or fracture. However, these alloys start to soften at anywhere between 1200 °C and 1400 °C. To render them capable of withstanding higher temperatures they are coated with a ceramic thermal barrier coating (TBC) comprising a material (often based on zirconia) with a low thermal conductivity. On its own, this TBC will not prevent the components from reaching the temperature of the inlet gases so the blades and other components must also be cooled internally by pumping air or steam through channels within them.

Air cooling requires hot air that is "stolen" from a later stage of the turbine. This has the effect of lowering overall efficiency but is simpler to design. In a combined cycle plant steam can be taken from the steam generator and used for cooling. This involves a smaller efficiency penalty but is more complex and can have an effect on overall plant flexibility. See Figure 4.4 for a diagram of gas turbine blade cooling.

While greater efficiency has been a primary goal for modern gas turbine manufacturers, another that has come to prominence since the beginning of the 21st century is flexibility. Flexibility is most important for large, high-efficiency combined cycle plants. Smaller, open cycle gas turbines can

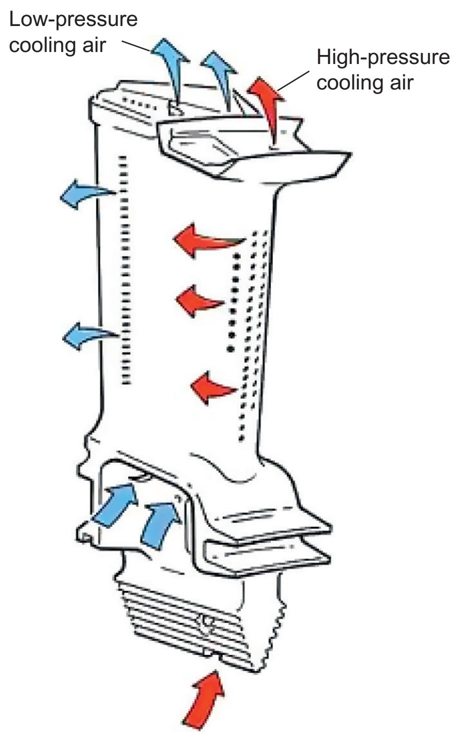

FIGURE 4.4 Gas turbine blade cooling.

generally be operated flexibly with little penalty. The goal of flexible operation is to be able to maintain good efficiency at part load as well as full load, and to be able to change load quickly. To achieve this, many involve sacrificing the ultimate level of efficiency. Flexible operation for combined cycle power plants will be examined in more detail later in the chapter.

ADVANCED GAS TURBINE CYCLES

The basic gas turbine configuration as just outlined can be modified in a number of ways to attempt to enhance the performance. Reheating, alluded to briefly already, involves adding a second combustion stage and splitting the turbine section into two parts. Intercooling is an analogous modification to the compressor in which this is divided into two stages, with cooling of the air in between them. Water vapor or steam injection into the inlet air before, within, or after the compressor can increase overall efficiency by creating a larger mass flow for less energy input. Finally, it is possible to capture some of the heat from the exhaust of a simple cycle gas turbine and use it to heat the compressed air from the compressor stage before it enters the combustion chamber, again leading to an improvement in efficiency. Some of these modifications can be utilized in gas turbines within combined cycle configurations, and others are only useful when the gas turbine is operating in an open cycle.

Reheating

In large steam turbine–based power plants it is traditional to split the turbine into separate sections: one handling high-pressure steam, one handling intermediate-pressure steam, and a third handling low-pressure steam. By splitting the turbine this way, efficiency gains can be made through matching the individual turbine sections to operate under a narrower range of steam pressures. Further, once the turbine has been split into separate sections, additional efficiency gains can be made by reheating the steam when it exits the high-pressure turbine (where it will have cooled) and before it enters the intermediate-pressure turbine. This is a common feature of the steam turbines used in coal-fired power plants.

The power turbine stage of a gas turbine can also be divided up in a similar way, though normally only two separate sections, called spools,[7] are used as shown in Figure 4.5. One of these is called the high-pressure spool and the second a low-pressure spool. Once the turbine has been split into spools it is possible to introduce a second combustion stage to reheat the air between the high-pressure and low-pressure spool of the power turbine. Reheating increases turbine efficiency in the same way as for a steam turbine by optimizing power

7. A *spool* is a set of turbine or compressor blades that rotate at the same speed.

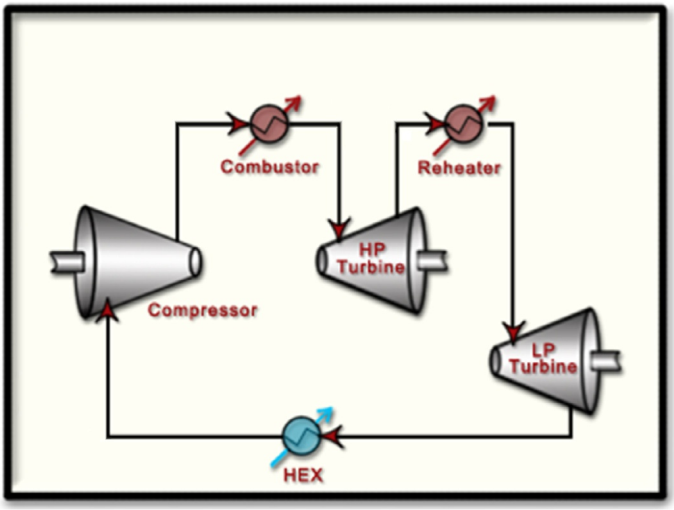

FIGURE 4.5 Gas turbine with reheat.

turbine and gas flow conditions. Further, it can allow the turbine to achieve higher efficiency with a lower turbine inlet temperature, making less demand on material performance.

Reheating is used by at least one major manufacturer[8] to achieve high efficiency in large gas turbines. These units also operate at a relatively high compression ratio of 30:1, similar to many aero-derivative gas turbines.

Intercooling

It is possible to go a stage further with a gas turbine by splitting the compressor into two sections: a low-pressure compressor section and a high-pressure compression section. As with the reheating of the air between the two spools of the turbine, it is possible to improve efficiency by cooling the air between the two spools of the compressor. (Compressing air generates heat, which raises the temperature of the air and hot air occupies a larger volume. Cooling it reduces the volume so the compressor actually has less work to do.) This is called intercooling.

Intercooling a high-performance aero-derivative gas turbine can boost its efficiency by around 5%, double its power output, and substantially reduce the cost per kilowatt of generating capacity.[9]

8. Alstom uses what it describes as sequential firing in its large gas turbines.
9. Arthur Cohn, *Humidified Gas Turbines*, presented at the Fourth Seminar on Combined Cycle Gas Turbines, British Institute of Mechanical Engineering, 1998.

Recuperation

A third strategy for improving the performance of a gas turbine is to use heat from the turbine exhaust to partially heat the compressed air from the compressor before it enters the combustion chamber. This process, referred to as recuperation, results in less fuel being needed to raise the air to the required turbine inlet temperature (Figure 4.6). In effect, the recuperated gas turbine cycle is allowing more energy to be captured from the air, increasing overall thermodynamic efficiency.

Recuperation uses heat in the gas turbine exhaust that would be captured in a combined cycle plant so this technique will not normally be applied in that configuration. However, it has been used successfully in a number of small gas turbines to provide higher efficiency. It is also a feature of many micro-turbines.

Mass Injection

A final strategy for increasing efficiency of a gas turbine is to inject water or water vapor into the air supplied to the turbine. Injection may take place at the compressor inlet, at stages through the compressor, or at the combustion chamber. In all cases the aim is to increase mass flow through the turbine.

The steam-injected gas turbine (STIG) cycle involves using a heat-recovery steam generator to produce steam from the exhaust gas of the gas turbine in much the same way as in a combined cycle station. This steam is then injected

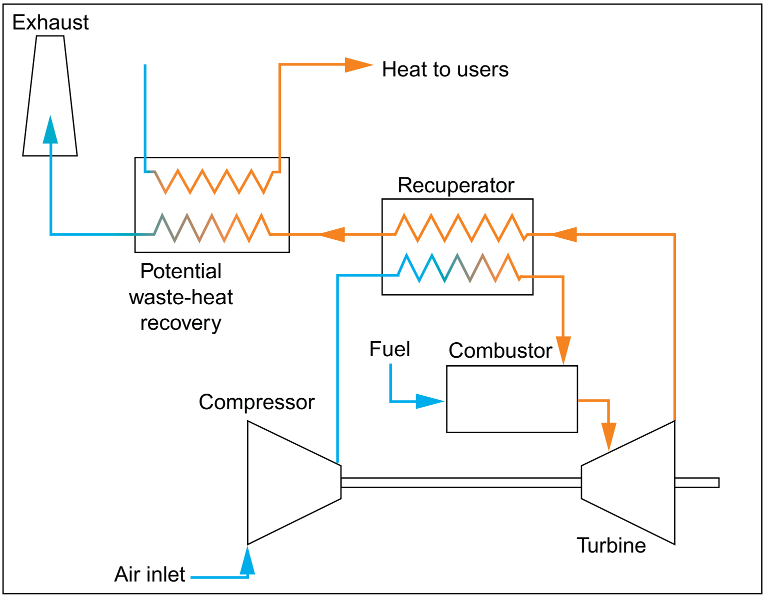

FIGURE 4.6 A gas turbine with recuperation.

into the combustion chamber of the turbine. This increases the mass flow into the turbine stages with a consequent increase in overall efficiency. An added advantage is that injection of steam into the combustion chamber can reduce Nitrogen oxide generation. A STIG cycle can increase overall gas turbine efficiency between 2% and 4%.

The humid-air turbine (HAT) cycle is a more complex cycle in which inlet air is intercooled part way through the compressor and at the exit of the compressor and then passed through a humidifier where it becomes almost saturated with water vapor before entering the combustion chamber. Recuperation is then employed to heat the humidified air and gain the highest possible efficiency in use of energy. For small gas turbines, this cycle is potentially more efficient than the equivalent combined cycle configuration.

Other HAT cycles are also possible. The cascaded HAT (CHAT) cycle introduces a reheating stage between the high-pressure spool and the low-pressure spool of the power turbine. Another variant called TOP Humid Air Cycle (TOPHAT) injects water into various stages of the compressor while the advanced HAT (AHAT) cycle injects water at the compressor inlet.

Various claims have been made for the HAT cycle and its variants with efficiencies as high as 55% in a large CHAT system. These suggest that in principle a HAT cycle turbine can approach the efficiency of a combined cycle power plant but without the complexity of a steam turbine and at lower inlet temperatures. However, the cycle cannot complete with the efficiency levels now being achieved in large combined cycle plants. Another disadvantage of HAT and CHAT cycle power units is that they release a considerable amount of water vapor into the environment. In situations where water is scarce it may be necessary to recover the water from the exhaust gas, adding to cost and complexity.

Along with the cycles described here there are a number of other variants of both the STIG and HAT cycles. Most have not been widely adopted commercially.

COMBINED CYCLE POWER PLANTS

The advanced cycles previously discussed offer the potential to increase gas turbine energy conversion efficiency. Where they have been applied, it has generally been for small generating systems. The most important adaptation of the gas turbine cycle, however, is the combined cycle power plant. This is a configuration that has been adopted and adapted by all gas turbine manufacturers, and it is capable of the highest energy conversion efficiency yet recorded for a large commercial fossil fuel–fired power station.

The efficiency of a gas turbine for electricity generation is always limited by the fact that the exhaust gases leave the turbine at a high temperature and therefore still contain a large amount of energy that has not been recovered. Some of the techniques described earlier attempt to use some of this heat energy. However, the most straightforward way of doing so is to add what is known as a bottoming cycle. This is an additional heat engine cycle operating on the low-temperature exhaust.[10] For a gas turbine, the best match for a bottoming cycle is a steam turbine.

To integrate the two, the exhaust from the gas turbine is fed into a specially designed heat-recovery steam generator that produces steam from the hot air. This steam is then used to drive a steam turbine generator that produces an additional amount of electricity (Figure 4.7). This basic design is capable of being interpreted in various ways. In some plants there will be a single steam turbine that is fed by steam generated from the exhausts of two or three gas turbines. In others, each gas turbine has its own steam turbine. In some cases these may be integrated on a single shaft with the gas turbine at one end, the steam turbine at the other end, and a generator in the middle driven by both. In this case it will usually be possible to decouple one or both of the turbines from the generator.

The development of gas turbine combined cycle plants since the late 1980s has led to a significant increase in energy conversion efficiency. In 1990, the best efficiency was around 50%. In 2011, a gas turbine combined cycle plant in Germany achieved 60.75% efficiency. Part of this improvement has been due to tight integration of all the combined cycle plant components to reduce heat loss. However, much of it has been due to an increase in the turbine inlet temperature, which has permitted greater thermodynamic efficiency. Inlet gas

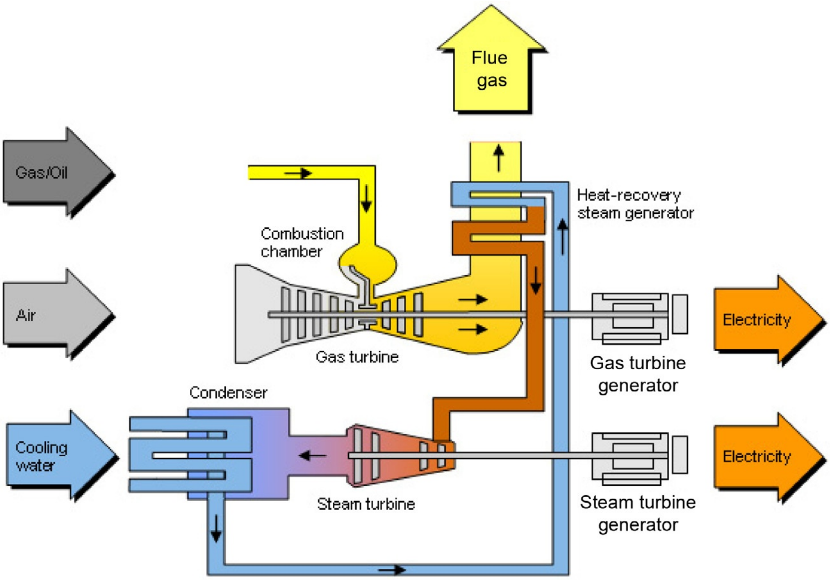

FIGURE 4.7 Schematic of a combined cycle power plant.

10. There is also a topping cycle that takes energy from the high-temperature, high-pressure gases.

temperatures have already reached 1600 °C in the hottest machines developed in Japan where a government program is aiming to make 1700 °C feasible, with a potential combined cycle efficiency of perhaps 65%.

The high efficiency of gas turbine power plants helps keep the cost of each unit of electricity down (though this depends critically on the price of natural gas). They are also attractive because they generate less carbon dioxide for each unit of electricity than other types of fossil fuel generation. In addition, the plants are cheap to build and can be erected relatively quickly because many of the components can be supplied to the power plant site already assembled.

Combined cycle plants can be built in a range of sizes for a few megawatts to several hundreds of megawatts, although sizes are limited by the unit sizes that the different manufacturers offer. While all have a part to play, it is the larger ones that have assumed the most importance in the market.

Originally these large combined cycle plants were conceived as base-load power generation units that would operate at full load for most of the time. The volatility of gas prices during the late 1990s and the first decade of the 21st century has often made this economically unviable and plants have frequently operated at capacity factors much lower than 100%. In the United States the typical capacity factor has been 40% or less in most recent years.

The duty cycle of the gas turbine plant has additionally been complicated by the increasing volumes of electricity entering grids in many parts of the world from renewable power stations exploiting wind or solar energy. The fluctuating output from these sources means that these grids require a secondary source of electricity to maintain the grid in balance. This is a role that combined cycle power plants are being adapted to fill.

To achieve flexibility, manufacturers are modifying the way their combined cycle plants operate. Flexible operation includes being able to generate at much less than 100% load while still maintaining good efficiency. It means being able to start up quickly and being able to change output quickly too, all without compromising either efficiency or emissions performance. Strategies to achieve this include keeping parts of the system continually warm by either heating them or by keeping the plant on-line at very low load (sometimes called parking) when demand is low—sometimes overnight—so that it can return to full output quickly. Various modifications to the way plants operate have allowed plants to change their outputs much more quickly than in the past.

Most large combined cycle plants have one to three gas turbines. However, some manufacturers are also exploring the option of using many more small gas turbines, each with its own steam generator and steam turbine, to provide greater flexibility. Each of the multiple units can then be brought into service as needed. Increasing or decreasing output is then managed by starting up or shutting down another unit so that operating units are at their optimum load for most of the time. This allows overall high efficiency to be maintained over a wide range of loads but at the expense of ultimate efficiency, because smaller gas turbines are generally less efficient than the largest machines.

MICRO-TURBINES

Micro-turbines are tiny gas turbines that can generate both electricity and heat. They vary in electrical output from around 25 kW to 250 kW. Units of between 250 kW and 500 kW are sometimes called mini-turbines. The two types are designed to be used in large domestic or small commercial environments where they can provide both forms of energy. As such, they are usually designated as distributed generators since they supply their power at the distribution level of the grid.

Depending on the source, micro-turbines are described as having evolved from turbochargers in automotive engines or auxiliary power units for aircraft. Whatever their origin, they are identical in their main components to conventional gas turbines with a compressor, a combustion chamber, and a power turbine (Figure 4.8). The main difference, apart from their size, is that a micro-turbine will only have one set of compressor blades and one set of turbine blades.

In most micro-turbines the generator is packaged with the turbine to provide a single unit that can be installed rapidly and with little preparation. As a consequence of their small size, micro-turbines and their generators operate at very high rotational speeds, typically between 40,000 rev/min and 120,000 rev/min. Generators rotating at these speeds cannot be connected directly to the grid. Instead they are connected through a solid-state interface that converts the high-frequency AC power produced by the generator to 50 Hz or 60 Hz as required to synchronize with the grid.

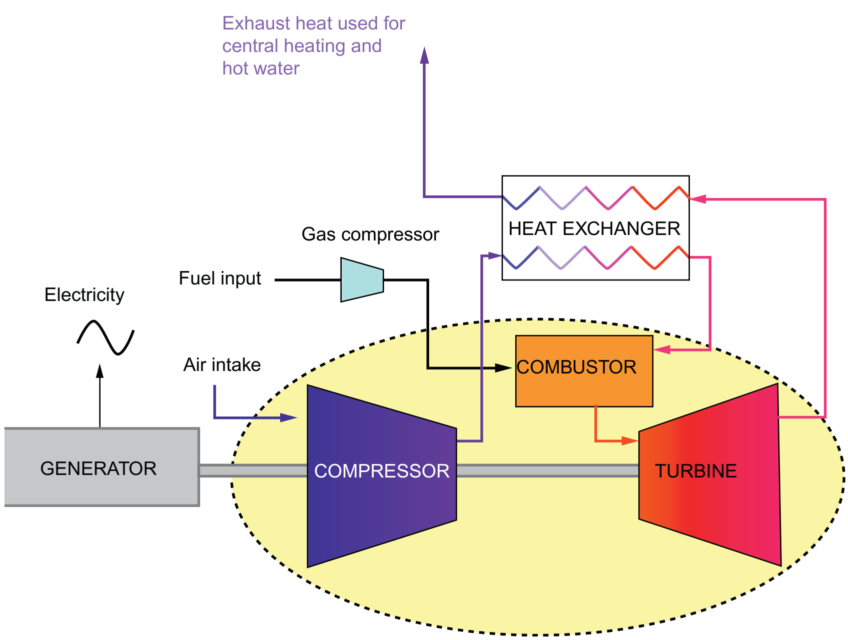

FIGURE 4.8 A schematic of a typical micro-turbine.

There are two general types of micro-turbine package available: simple and recuperated. The simple micro-turbines tend to be the more robust but their energy conversion efficiency is low at around 15%. However, this leaves a significant amount of waste heat in the turbine exhaust for supplying hot water in a cogeneration application. Recuperated micro-turbines use waste heat to heat compressed air between the compressor stage and the combustion chamber in exactly the way as industrial and aero-derivative gas turbines. Recuperation on a micro-turbine can increase overall conversion efficiency between 20% and 30%. With both types, waste heat capture for hot water can raise overall cogeneration efficiency to 85%.

Most micro-turbines are designed to burn natural gas. A packaged system will comprise the micro-turbine generator, a gas compressor to provide natural gas at the pressure required by the turbine combustion chamber, a solid-state grid interface, and, in many cases, a waste heat-recovery system. They have very low emissions and can be installed in domestic and commercial environments. Micro-turbines are more expensive than similarly sized reciprocating engines, but research is continuing with the aim of achieving much higher energy-to-electrical conversion efficiencies.

EMISSION CONTROL FOR GAS TURBINE POWER PLANTS

The combustion of natural gas, the main fuel for gas turbines, has a relatively low environmental impact when compared to its main fossil fuel competitor, coal. After it has been cleaned it contains relatively little if any hydrogen sulfide and no heavy metals. Depending on the precise composition of the natural gas and on the gas turbine combustion system in which it is burned, its combustion will generate some carbon monoxide and also some particulate material, both resulting from incomplete combustion of components of the fuel. Aside from these, the main atmospheric pollutant of concern with gas turbine power generation is nitrogen oxide generated by oxidation of nitrogen in the combustor. Levels of carbon monoxide, particulates, and nitrogen oxide can peak during startup and under part-load operation.

Gas turbines also generate carbon dioxide from the combustion of the hydrocarbon fuel they burn. The quantity generated for each unit of electricity is much less than would be released from a coal-fired power station, but even so, large combined cycle power plants are major carbon dioxide emitters.

Nitrogen Oxide

Nitrogen oxide can be produced in gas turbines from two different sources. Small amounts of nitrogen can be found bound in the fuel itself as part of the complex mixture of hydrocarbon-based materials that can make up natural gas. This nitrogen, when it exists, will be converted into nitrogen oxide during combustion. However, the main source of nitrogen oxide is from nitrogen gas, either from the air in which the gas is burned or actually contained in small amounts with the fuel.

Most nitrogen oxide production is the result of the oxidation of nitrogen gas by oxygen in the air at the high temperatures that are reached in the combustor of the gas turbine. With power turbine inlet temperatures as high as 1600 °C or greater in large high-efficiency gas turbine combined cycle systems, this presents a significant problem. The solution has been to develop low nitrogen oxide burners that control the combustion in such a way as to limit nitrogen oxide production.

One successful technique that has been applied in a range of gas turbine combustors is water or steam injection. This serves to lower the combustion temperature and thereby reduce nitrogen oxide production. However, as temperatures have risen higher, these have been replaced by dry low nitrogen oxide burners that achieve the same end without water injection.

The simplest way to burn natural gas in air is to pump the gas through a fine nozzle where it is ignited as it enters the atmosphere. With unlimited amounts or air, and therefore of oxygen available, this diffusion or spray combustion results in a stable flame within which optimum combustion conditions will ultimately be reached, leading to complete combustion and limited formation of carbon monoxide and particulate material. However, this unrestrained combustion can also lead to high levels of nitrogen oxide.

The main alternative and the technique that is now used in most dry low nitrogen oxide burners is called premixed combustion. This involves premixing the natural gas with a carefully controlled quantity of air before it enters the combustion chamber, such that there is just sufficient oxygen to react with the combustible gas but none left to react with nitrogen. Controlling combustion under these conditions is much more difficult. If there is too little air then combustion is incomplete and will result in high concentrations of carbon monoxide and unburned material, leading to harmful emissions and lower efficiency because all the energy within the fuel is not utilized. On the other hand, if the amount of air becomes too high, then nitrogen oxide levels quickly rise.

Therefore, premixed combustion must be carefully controlled. When it is, modern low nitrogen oxide burners are capable of reducing the emissions of nitrogen oxide to between 15 ppm to 25 ppm for large industrial gas turbines and as low as 9 ppm for some smaller gas turbines. Nitrogen oxide emission standards are typically 25 ppm in many parts of the world and manufacturers aim to meet these standards without further control being necessary. However, standards are tightening and some areas are introducing emission limits of 15 ppm or 10 ppm for gas turbines. These are much harder for large combined cycle plants to meet without further nitrogen oxide reduction systems.

Where low nitrogen oxide burner technology cannot meet the required emission standard, manufacturers have to turn to an alternative post-combustion technology. The most common of these is selective catalytic reduction (SCR). SCR involves injecting a reducing gas, usually ammonia or, sometimes, urea, into the exhaust gases after they exit the power turbine and then passing the mixture over a metal catalyst that catalyzes the reaction between the ammonia, or urea, and nitrogen oxide to produce nitrogen and water vapor. In a combined cycle power plant the SCR unit is often placed within the heat-recovery steam generator as shown in Figure 4.9.

Catalysts for SCR are normally either a metal or metal oxide on a ceramic carrier. Various metals have been used, including vanadium, molybdenum, and tungsten, as well as platinum, though the latter is mostly found in the system used in car exhausts. The type of catalyst will depend on the temperature of the exhaust gases. Metal oxides are more temperature resistant than metals.

In principle, an SCR-based nitrogen oxide reduction system can remove 95–99% of the nitrogen oxide from the exhaust gases. However, the system becomes more difficult to control when reduction levels exceed 80% because the reaction does not proceed as smoothly, putting greater demands on the catalyst and often leading to higher levels of ammonia or urea passing through the system and being released into the atmosphere. This is referred to as ammonia slip.

Therefore, the optimum solution is a balance between low nitrogen oxide burners and SCR such that both operate within the best range of efficiency. With both, it is possible to reduce nitrogen oxide emission levels below 10 ppm and often much lower than this. Targets of 4–5 ppm appear to be achievable with modern systems.

Carbon Monoxide

Carbon monoxide can be generated in gas turbines as a result of incomplete combustion of the natural gas fuel. Emissions of carbon monoxide are controlled in the same way as those of nitrogen oxide with similar limits of

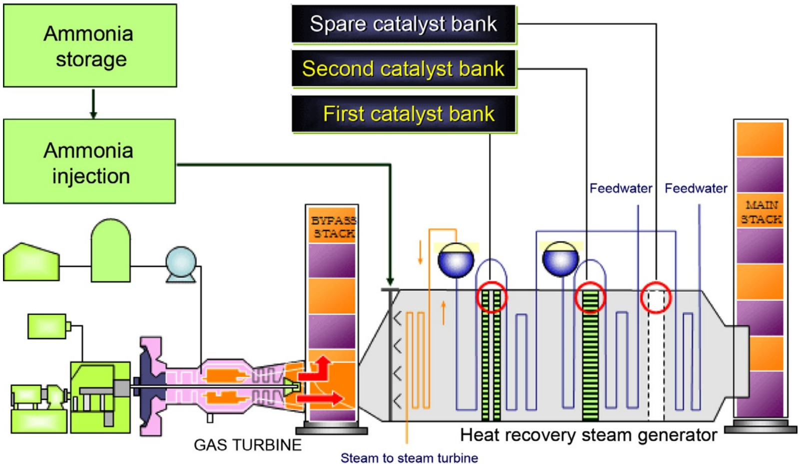

FIGURE 4.9 Selective catalytic reduction in a combined cycle power plant heat-recovery steam generator.

10–25 ppm in operation. Low nitrogen oxide burner technology can lead to levels of carbon monoxide higher than this, under which circumstances some system of emission control is needed. This is usually in the form of an oxidation catalyst that catalyzes the conversion of carbon monoxide into carbon dioxide. This may be a separate catalytic unit, but in combined cycle plants it is often incorporated into the heat-recovery steam generator too.

Carbon Dioxide

The combustion of natural gas generates significant quantities of carbon dioxide. Although the rate of production is much lower than in a coal-fired power plant, the gross level of production can be large. In the United States in 2009, for example, gas turbine–based power plants accounted for 23% of total power generation and 13% of carbon dioxide emissions.

For each kilowatt of energy contained in natural gas, its combustion will produce 0.23 kg of carbon dioxide. For coal the equivalent figure is 0.37 kg/kWh. The actual emissions per unit of electricity produced will depend on the efficiency of electricity production, but a combined cycle plant operating at 55–60% efficiency will emit relatively less than a coal plant at 38–45% efficiency.

Combined cycle plants are considerably less carbon intensive than coal-fired plants, and this has been one of the driving forces behind a switch to more natural gas–fired power generation in the developed world. Even so, when carbon capture from power plants becomes necessary, it is unlikely that combined cycle units will be exempt.

The main methods of reducing or eliminating carbon emissions from gas turbine–based power plants are exactly the same as those for coal-fired power plants that were discussed in Chapter 3: post-combustion capture, pre-combustion capture, and oxy-fuel combustion. The concentration of carbon dioxide in the flue gases from a typical combined cycle power plant will be 3–4%. This is much lower than in the flue gases of a coal-fired power plant and makes post-combustion capture relatively more difficult. On the other hand, the low concentration means that less needs to be removed.

Post-combustion capture using an ammonia or monoethanolamine absorbent in a spray tower is likely to be the most effective form of post-combustion capture for combined cycle plants. It is also an efficient means of retrofitting capture technology. The problem of low carbon dioxide concentration in the exhaust gases can be mitigated by recycling exhaust gases back to the gas turbine inlet. This has the effect of increasing carbon dioxide concentration but also of reducing the available oxygen entering the combustion chamber. Exhaust gas recycling of up to 35% appears to be effective without adversely affecting the combustion conditions. With exhaust gas recycling, post-combustion capture should be able to remove 90% of the carbon dioxide generated. Overall energy conversion efficiency is reduced by perhaps 6–8%.

Pre-combustion capture is essentially analogous to the coal gasification technology discussed in Chapter 3, producing hydrogen fuel gas that is then burned in a gas turbine combined cycle power plant. In the case of natural gas the process is known as reforming and is widely carried out industrially to make hydrogen. As with post-combustion capture, the capture efficiency is likely to be around 90%, but the overall generation efficiency is likely to fall more, perhaps by 10%.

Oxy-fuel combustion involves burning the natural gas in pure oxygen instead of air. This requires an oxygen separation plant but leaves an exhaust gas that is rich in carbon dioxide and containing little nitrogen, making it much easier to isolate. Capture efficiency may be as high as 99%. As was the case in a coal-based oxy-fuel plant, the combustion temperature when natural gas is burned in oxygen is far higher than in air—too high for the materials currently available to withstand. To counter this, carbon dioxide–rich exhaust gases are fed back to the gas turbine inlet to dilute the oxygen, reducing the combustion temperature to a level similar to that for air combustion. Overall energy conversion efficiency is likely to fall by up to 8–9% compared to a gas turbine plant without capture.

COST OF GAS TURBINE–BASED POWER GENERATION

The capital cost of a gas turbine–based power plant depends primarily on the cost of the gas turbine. This is a high-technology component that is manufactured by a small number of companies, mostly based in developed countries. However, even the largest gas turbines are off-the-shelf components that can be supplied virtually ready to operate. Competition between different manufacturers is stiff, and in many cases the key negotiating tool is price. As a consequence, gas turbines tend to be extremely competitively priced and they represent one of the least capital cost forms of power generation, particularly in the shape of a large combined cycle power plant.

Against that, the fact that gas turbines are high-technology components means that their cost will always be affected by shifts in commodity prices for the materials, especially the metals, from which they are built. This will be magnified in a combined cycle plant as a result of the cost of the heat-recovery steam generator, made primarily from steel and the steam turbine, again built mostly from steel. However, the fact that major components of a combined cycle plant can be delivered ready-assembled reduces labor costs.

Table 4.3 contains costs for gas turbine–based power plants in the U.S. market, based on analysis by the U.S. Energy Information Administration (EIA). While costs in other parts of the world will differ, the fact that these units are traded, internationally, means that prices are likely to be similar, whatever market is being considered. Based on this U.S. EIA analysis, which considers costs for plants in year 2010 U.S. dollars and where the order was placed in 2011, the cost of a conventional natural gas–fired combined cycle plant is

TABLE 4.3 Capital Cost of Gas Turbine–based Power Generation

Type of Gas Turbine Plant	Plant Capacity (MW)	Capital Cost ($/kW)
Conventional gas turbine combined cycle plant	540	931
Advanced gas turbine combined cycle plant	400	929
Advanced gas turbine combined cycle plant with carbon capture and storage	340	1834
Conventional open cycle gas turbine plant	85	927
Advanced open cycle gas turbine plant	210	634

Note: *The figures used in the table are base overnight costs in 2010.*

Source: *Assumptions to the Annual Energy Outlook 2012,* U.S. Energy Information Administration.

$931/kW. An advanced combined cycle plant with the most efficient gas turbines available would actually be slightly cheaper at $921/kW. Typical plant sizes for these two types of plant are shown in Table 4.3.

Carbon capture and storage is expected to increase the cost of fossil fuel power generation considerably. In Table 4.3, the cost for an advanced gas turbine combined cycle plant with carbon capture and storage is estimated to be $1834/kW, or slightly under twice the cost of either type of plant without carbon capture and storage. A conventional open cycle gas turbine plant, likely to be used primarily for peak power and grid support services, costs $927/kW or virtually the same as the combined cycle plants without carbon capture and storage. This plant would, however, operate with significantly lower efficiency than the combined cycle plant, as explained before. Meanwhile, an advanced gas turbine in an open cycle configuration is expected to cost $634/kW. The two types of open cycle plant are expected to have a lead time from order to entering service of two years. For the combined cycle plants the lead time would be three years.

Table 4.4 shows some future costs for combined cycle power plants burning natural gas based on a report from the U.S. Electric Power Research Institute (EPRI). Again, the figures apply to the U.S. market but should be broadly indicative of prices elsewhere. EPRI calculated that a combined cycle power plant without carbon capture entering service in 2015 would have a capital cost of $1060–1150/kW. EPRI expects the cost to be broadly unchanged in 2025. Meanwhile a similar plant, but with carbon capture, is predicted to cost between $1600/kW and $1900/kW.

The low capital cost of gas turbine–based power generation is one of the reasons why they have been favored by private sector companies operating in the liberalized electricity markets established around the world. The cost

TABLE 4.4 Capital Costs in 2015 and 2025 for U.S. Gas-fired Combined Cycle Plants

Plant	2015 ($/kW) without Carbon Dioxide Capture	2025 ($/kW) Carbon Capture and Storage
Natural gas combined cycle plant	1060–1150	1060–1150
Natural gas combined cycle plant with carbon capture	n/a	1600–1900

Source: Program on Technology Innovation: Integrated Generation Technology Options, Technical Update, June 2011, Electric Power Research Institute, 2011. (EPRI reference 1022782.)

remains low compared to most other technologies even when carbon capture and storage is included. However, that does not necessarily translate into a low cost for electricity generated by these plants.

The main problem with natural gas–based power generation is the cost of gas. When this is low, combined cycle power plants are extremely competitive sources of electric power. However, when prices rise, and in global energy markets the price can rise extremely rapidly, they can become far from economical. The uncertainty attached to the volatility in gas prices can also have a damaging effect on the economic growth in countries that rely heavily on gas-fired generation.

The IEA predictions of a golden age of gas, if they come to pass, may lead to a loosening of gas supplies, a lowering of average prices, and a reduction in volatility. If so, then gas-fired generation is likely to thrive well into the 21st century.

Chapter 5

Piston Engine–Based Power Plants

Piston engines or reciprocating engines (the two terms are often used interchangeably) are by a wide margin the largest group of thermodynamic heat engines in use around the world. Their applications range from model airplanes to lawn mowers, and they include all the automotive power plants found in motorcycles, cars, trucks, and many other sorts of heavy machinery. They also power locomotives, ships, and many small aircraft, and they provide stationary electrical power and combined heat and power to numerous sites across the globe.

The number in use is enormous; the United States alone produces more than 35 million each year. Engines vary in size from less than 1 kW (model engines can be a few watts) to 65,000 kW. They can burn a wide range of fuels including natural gas, biogas, LPG, gasoline, diesel, biodiesel, heavy fuel oil, and even coal. They are manufactured all over the globe and there is a large global base of expertise in their maintenance and repair.

In line with the wide range of engines available, the power generation applications of piston engines are enormously varied. Small units can be used for standby power or for combined heat and power in homes and offices. Larger standby units are often used in situations where a continuous supply of power is critical, such as in hospitals or to support highly sensitive computer installations like for air traffic control and the many server farms around the world. Commercial and industrial facilities use medium-size piston engine–based combined heat and power units for base-load power generation. Large engines, meanwhile, can be used for base-load, grid-connected power generation, while smaller units form one of the main sources of base-load power to isolated communities with no access to an electricity grid.

The piston engines used for power generation are almost exclusively derived from similar engines designed for motive applications. Smaller units are normally based on car or truck engines, while the larger engines are based on locomotive or marine engines. Performance of these engines vary. The small engines are usually cheap because they are mass produced, but they have relatively low efficiencies and short lives. Larger engines tend to be more expensive, but they will operate for much longer. Large, megawatt-scale engines are

Power Generation Technologies

among the most efficient prime movers available,[1] with simple cycle efficiencies approaching 50%.

There are two main types of reciprocating engines: the spark-ignition engine and the compression or diesel engine. The latter was traditionally the most popular for power generation applications because of its higher efficiency. However, it also produces high levels of atmospheric pollution, particularly nitrogen oxide. As a consequence, spark-ignition engines burning natural gas have become the more popular units for power generation, at least within industrialized nations. A third type of piston engine, called the Stirling engine, is also being developed for some specialized power generation applications. This engine is novel because the heat energy used to drive it is applied outside the sealed piston chamber.

INTERNAL COMBUSTION ENGINES

The common feature of all reciprocating engines is a cylindrical chamber housing a piston. In its most basic form the engine comprises a single cylinder sealed at one end and open at the other end. A cylindrical-shaped disk of metal, the piston, is designed to fit closely within the cylinder to seal the open end, and this piston can move backwards and forwards easily within the cylinder. This it does in response to the pressure changes in the gas contained within the cylinder at various stages in the engine cycle. Meanwhile, air and fuel can be admitted into the sealed end of the cylinder, and exhaust gases removed after combustion through moveable valves. The outside of the piston is attached via a hinged lever to the main drive shaft of the engine. Movement of the piston in and out of the cylinder is converted via this hinged linkage into rotary motion of the main shaft and this rotary motion is used to derive power or electricity.

To convert chemical energy contained within a fuel into mechanical energy, an air–fuel mixture is admitted into the cylinder of the engine and ignited, causing a controlled explosion and a high-pressure impulse from which the power of the engine is derived. The pressure impulse forces the gases in the cylinder to expand, pushing the piston to the limit of its movement. This controlled explosion gives the engines its common name: the internal combustion engine.

There are two main types of internal combustion engine: the spark ignition engine and the diesel or compression-ignition engine, each defined by the way in which fuel is admitted into the engine cylinder and how ignition of the air–fuel mixture is initiated. The two types of engine can be further subdivided according the engine cycle they employ. Two are common: the two-stroke cycle and the four-stroke cycle. Internal combustion engines account for virtually all the reciprocating engines in use around the world.

1. Slow-speed engines are the most efficient engines for converting fuel energy via heat into rotary motion to generate electricity. Fuel cells, which turn chemical energy directly into electrical energy, can be more efficient.

Of the two common engine cycles, the four-stroke cycle is the most commonly used, particularly for small- and medium-size engines. A two-stroke cycle is employed is some very small engines because it can provide a high power-to-weight ratio. It is also often used in very large engines since it is capable of high efficiency and tolerating very poor fuels.

Both cycles are broadly similar in concept to the principle employed in the gas-turbine engine described in Chapter 4. Air is compressed and then fuel added and ignited, generating a hot, high-pressure gas from which energy can be extracted in the form of mechanical work. However, whereas the gas-turbine engine is continuously compressing air, mixing it with fuel and igniting the mixture, the piston engine carries these processes out sequentially.

Engine Cycles

The internal combustion engine is a thermodynamic heat engine, and as such belongs to the same category as steam turbines and gas turbines. However, the physical nature of the reciprocating engine is very different to that of the turbines. The reciprocating engine principle was developed in the latter half of the 19th century, though some primitive engines were in existence before that. Nikolaus Otto is generally credited with building the first four-stroke internal combustion engine in 1876. In doing so, he established the principle still in use today.

The Otto cycle engine employs a spark to ignite a mixture of air and, traditionally, gasoline[2] compressed by the piston within the engine cylinder. This spark ignition causes an explosive release of heat energy that increases the gas pressure in the cylinder, forcing the piston outwards as the gas tries to expand. This explosion is the source of power, its force on the piston turning the crankshaft to generate rotary motion.

The Otto cycle was modified by Rudolph Diesel in the 1890s. In his version, air is compressed in a cylinder by a piston to such a high pressure that its temperature rises above the ignition point of the fuel, which is then introduced into the chamber and ignites spontaneously without the need for a spark. This represents a simplification of the Otto cycle but is not without its complications, particularly from an emissions perspective.

In a four-stroke engine each piston of the engine—and there can be a large number depending on the particular engine type and application—is equipped with at least two valves: one to admit air or an air–fuel mixture, and a second to exhaust spent gases after ignition. The opening and closing of these valves is mechanically synchronized with the movement of the piston backwards and forwards.

The four-stroke cycle derives its name from the four identifiable movements of the piston in the chamber, two of expansion and two of compression, for each

2. Otto's engine probably burned powdered coal but gasoline soon became the preferred fuel.

full power cycle. Starting with the piston at the top of its chamber, and the chamber empty, the first stroke is an intake stroke in which either air (diesel cycle) or a fuel and air mixture (Otto cycle), is drawn into the piston chamber by movement of the piston to expand the volume of the enclosed space (Figure 5.1) and the air or fuel–air valve is open. This valve closes at the end of the first stroke. The second stroke is a compression stroke during which the gases in the cylinder are compressed by the piston returning toward the top of its chamber. In the case of the Otto cycle, a spark ignites the fuel–air mixture at the top of this second stroke, creating an explosive expansion of the compressed mixture that forces the piston down again. This is the power cycle. In the diesel cycle fuel is introduced through a separate nozzle close to the top of the compression stroke, igniting spontaneously in the hot gas with the same effect. After the power stroke, the fourth stroke is the exhaust stroke during which the exhaust gases are forced out of the piston chamber through the second valve, which is now open. This closes at the end of the fourth stroke and the cycle begins again. In both spark-ignition and diesel engines a large flywheel attached to the crankshaft stores angular momentum generated by the power stroke and this provides sufficient momentum to carry the crankshaft and piston through the three other stokes required for each cycle.

The shaft of an engine that is fitted with a single piston and cylinder will receive a power impulse once every two rotations, leading to a relatively uneven transfer of power. However, if the engine has multiple cylinders, the cycle of each can be staggered relative to the others so that they deliver their power sequentially, leading to a much more even rotational motion. For a four-stroke engine it is normal for four (or a multiple of four) pistons to be attached to the crankshaft, with one of each set of four timed to produce a power stroke while the other three move through different stages of their cycles. The introduction of

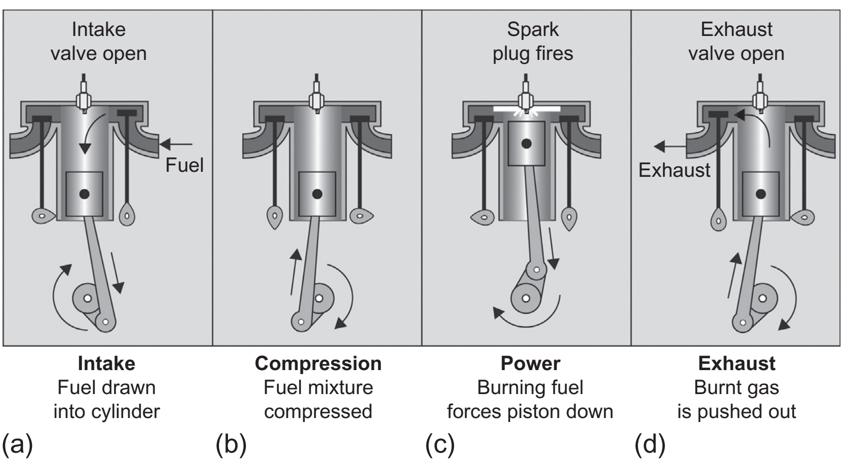

FIGURE 5.1 The strokes of a four-stroke cycle.

fuel and air and the removal of the exhaust gases are then controlled by valves that are mechanically timed to coincide with the various stages of the cycle on each cylinder.

In a two-stroke engine, intake and exhaust strokes are not separate. Instead, fuel is forced into the piston chamber (intake) from an opening in the crankshaft casing toward the end of the power stroke, pushing out the exhaust gases through a second opening at the top of the chamber. The same stroke then compresses the fuel–air mixture within the cylinder, the exhaust port having been closed by movement of the piston. This exhaust–compression stroke is then followed by ignition of the fuel and a power–expansion stroke, then a repeat of the cycle (Figure 5.2).

Two-stroke engines are simpler than four-stroke engines because they do not require valves. In addition, power is delivered by each cylinder at each revolution of the shaft rather than once every two revolutions. This means they can deliver more power for a similar weight.

Engine Size and Engine Speed

The speed at which a piston engine operates will usually depend on its size. In general, small units operate at the highest shaft rotational speed and large units at the lowest shaft speed. In addition, in most situations a piston engine–based power unit will have to be synchronized to an electricity grid operating at 50 Hz or 60 Hz, so the engine speed will also be determined by one or the other of these rates. So, for example, a 50 Hz high-speed engine will typically operate at 1000 rpm, 1500 rpm, or 3000 rpm, while the equivalent 60 Hz machine will operate at 1200 rpm, 1800 rpm, or 3600 rpm. These speeds allow the generators attached to the engines to synchronize with the grid operating frequency.

Engines are usually classified according to speed into one of three groups, high speed, medium speed and slow speed engines. High speed engines are the smallest and operate up to 3600 rpm. The largest slow speed engines may run as slow as 50 rpm. Typical speed and power ranges for each type of engine are shown in Table 5.1.

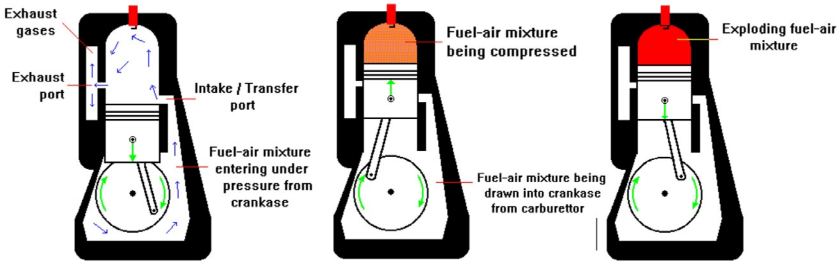

FIGURE 5.2 The two-stroke cycle.

TABLE 5.1 Piston Engine Classification by Size and Speed

	Engine Size	Engine Speed (rpm)
High speed	1 kW–8.5 MW	1000–3600
Medium speed	1 MW–35 MW	275–1000
Slow speed	2 MW–65 MW	50–275

Source: *Technology Characterization: Reciprocating Engines,* U.S. Environmental Protection Agency, 2002.

Engine performance varies with speed. High-speed engines provide the greatest power output as a function of cylinder size, and therefore the greatest power density. However, the larger, slower engines are more efficient and last longer. Thus, the choice of engine will depend very much on the application for which it is intended. Large and slow- or medium-speed engines are generally more suited to base-load generation, but it may be more cost effective to employ high-speed engines for backup service where the engines will not be required to operate for many hours each year.

In addition to standby service or continuous output base-load operation, piston engine power plants are good at load following. Internal combustion engines operate well under part-load conditions. For a gas-fired spark-ignition engine, output at 50% load is roughly 8–10% lower than at full load. The diesel engine performs even better, with output barely changing when load drops from 100% to 50%.

SPARK-IGNITION ENGINES

Spark-ignition engines are capable of burning a variety of fuels, including gasoline, propane, biogas, and landfill gas. In practice, however, many of these engines burn natural gas when used for power generation applications because of the lower emissions. These power-generating units are generally four-stroke engines and they are available in sizes from less than 1 kW up to around 6.5 MW (Table 5.2).

The spark-ignition engine uses a spark plug to ignite the fuel–air mixture that is admitted to each cylinder of the engine. In the simplest case this spark plug is located in the top of the cylinder and directly ignites the mixture within the cylinder. The composition of the fuel–air mixture in the cylinder may be close to the stoichiometric ratio required for complete combustion of the fuel, but more often it will contain a significant excess of air. More technically complex engines can use a preignition chamber in which a small amount of a fuel–air mixture rich in fuel is admitted and ignited. This preignition then spreads into the main cylinder where a fuel–air mixture containing a much greater proportion of air is ignited.

TABLE 5.2 Four-stroke Engine Performance Parameters

	Diesel Engine	Spark-ignition Engine
Typical size range	1 kW–65 MW	1 kW–6.5 MW
Efficiency	20–48%	28–42%
Compression ratio	14:1 to 25:1	8:1 to 12:1

In common with all thermodynamic heat engines, the efficiency that a reciprocating engine can achieve increases with the temperature of the working fluid: air. For a spark-ignition engine the highest cylinder temperature is reached when the air-to-fuel ratio is around 16:1, the ratio at which a stoichiometric amount of oxygen is available to react with the fuel. An engine that operates with this air–fuel mixture is described as a rich-burn engine. A rich mixture leads to the highest temperature but it also leads to the greatest formation of nitrogen oxide, as well as significant amounts of carbon monoxide and unburned hydrocarbon particles as a result of incomplete combustion of some of the fuel. Under most circumstances, therefore, engines operating on a rich mixture will require emission control systems to limit the release of these potential pollutants.

If engine emissions are to be reduced during combustion, then the combustion temperature must be lowered and a greater amount of oxygen introduced to allow complete combustion of the fuel. Such engines are described as lean-burn engines and can operate with an air-to-fuel ratio between 20:1 and 50:1, significantly higher than in the rich-burn engine. The greater proportion of air lowers the overall combustion temperature (there will be less fuel entering the combustion chamber in the lean mixture), reducing the production of nitrogen oxide from nitrogen in the air, and provides the conditions for much more complete combustion of the fuel. This will reduce the amounts of carbon monoxide and unburned hydrocarbons in the exhaust gases. Against this, the lower temperature reduces overall efficiency. Lean-burn engines achieve a typical efficiency of only 28% (LHV),[3] compared to up to 42% (LHV) for a rich-burn engine. An engine tuned for maximum efficiency will produce roughly twice as much nitrogen oxide as one tuned for low emissions. Typical nitrogen oxide emission levels for spark-ignition engines are 45 ppmV to 150 ppmV.

The compression ratio of a spark-ignition engine (the amount by which the air–fuel mixture is compressed within the cylinder) is normally limited to a

3. The energy content of a fuel may be expressed as either the higher heating value (HHV) or the lower heating value (LHV). The higher heating value represents the energy released when the fuel is burned and all the products of the combustion process are then cooled to 25 °C. This energy then includes the latent heat of vaporization released when any water produced by combustion of, for example, natural gas, is condensed to room temperature. The LHV does not include this latent heat and is therefore around 10% lower than the HHV in the case of natural gas.

maximum between 9:1 and 12:1 to prevent the mixture from becoming too hot and spontaneously igniting, a process known as knocking. Lean natural gas–air mixtures have a much higher resistance to knocking than stoichiometic mixtures and can tolerate higher compression ratios than gasoline.

DIESEL ENGINES (COMPRESSION ENGINES)

Whereas a spark-ignition engine draws a fuel–air mixture into the cylinder during the first stroke of the four-stroke cycle, in a diesel engine only air is admitted at this stage (Figure 5.3). This air is then compressed much more highly during the succeeding compression stroke than would be the case with a spark-ignition engine, with compression ratios of up to 25:1 typical. When air is compressed adiabatically in this way (just like the air being compressed in a bicycle pump), the compression generates heat so that the air becomes much hotter. In the diesel engine the compression ratio is chosen so that the air temperature rises so high that it is above the ignition temperature of the fuel, in this case diesel. Fuel is admitted into the chamber under pressure toward the top of the compression stroke, when it then ignites spontaneously.

The temperature inside the cylinder of a diesel engine during ignition rises much higher than the temperature in a spark-ignition engine. As a consequence the production of nitrogen oxide is much higher. Typical levels are 450 ppmV to 1800 ppmV, or 10 times higher than for the equivalent spark-ignition engine (see Table 5.3 later, but note that as shown in Table 5.4 that they can rise even higher). Diesel engines, therefore, require extensive emission control systems if they are to comply with air-quality regulations, particularly when the units are operating in an urban environment. Against this, diesel engines are capable of burning biodiesel, a carbon dioxide–emission neutral fuel. This can be attractive in some situations.

The efficiency of the diesel engine ranges from 30% (HHV) for small engines to 48% (HHV) for the largest engines. The highest efficiency yet achieved is 52% for a 67 MW marine diesel engine. Higher temperatures could conceivably produce higher efficiency, but as with other types of

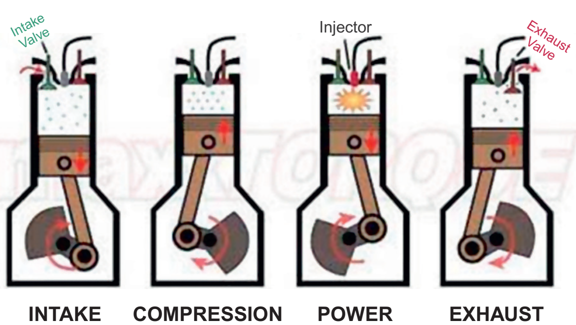

FIGURE 5.3 Four-stroke diesel engine cycle.

TABLE 5.3 Emissions of Nitrogen Oxide from Internal Combustion Engines

	Emissions (ppmV)	Emissions (g/kWh)
High-speed and medium-speed diesel engine	450–1800	7–20
Spark-ignition natural gas engine	45–150	1–3

Source: *Technology Characterization: Reciprocating Engines,* U.S. Environmental Protection Agency, 2008.

TABLE 5.4 Range of Emissions from Diesel Engines

Emissions	Emission Range
Nitrogen oxide	50–2500 ppmV
Carbon monoxide	5–1500 ppmV
Particulate matter	0.1–0.25 g/m^3
VOCs	20–400 ppmV
Sulfur dioxide	10–150 ppmV

Source: Nett Technologies.

thermodynamic engines, materials will be the limiting factor. Diesel engines can be built to larger sizes than spark-ignition engines, with high-speed machines available in sizes up to 4 MW and slow-speed diesels up to 65 MW. Large, slow-speed engines can have enormous cylinders. For example, a nine-cylinder, 24 MW engine used in a power station in Macau has cylinders with a diameter of 800 mm.

Diesel engines can burn a range of diesel fuels including both oil-derived fuels and biofuels. Smaller, high-speed engines normally use high-quality distillate but the large, slow-speed engines can burn very low-quality heavy fuel oils, which require a much longer combustion time to burn completely. These fuels tend to be dirty and plants burning them usually require additional emission mitigation measures.

DUAL-FUEL ENGINES

The large disparity in efficiency between a spark-ignition engine and a diesel engine has prompted engine developers to search for a way on achieving the efficiency of a diesel engine in a natural gas–fired spark-ignition engine. This

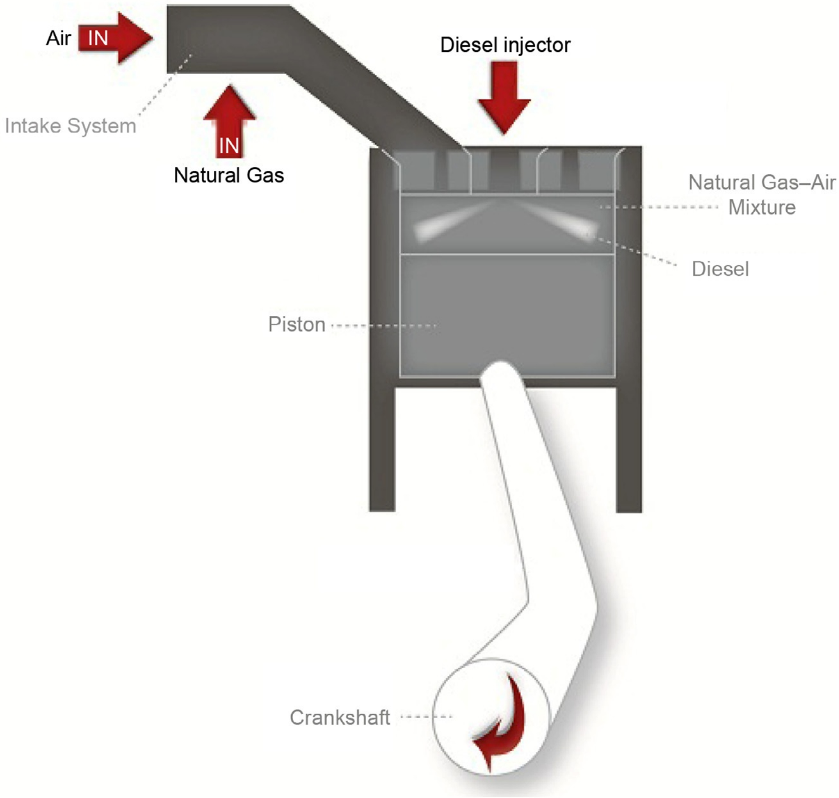

FIGURE 5.4 Schematic of a dual-fuel engine.

is the origin of the dual-fuel engine, which has been the most successful of these hybrids (Figure 5.4).

A dual-fuel engine is an engine designed to burn predominantly natural gas but with a small percentage of diesel as a pilot fuel to start ignition. The engines operate on a cross between the diesel and the Otto cycles. In operation, a natural gas–air mixture is admitted to the cylinder during the intake stroke, then compressed during the compression stroke. At the top of the compression stoke the pilot diesel fuel is admitted and ignites spontaneously, igniting the gas–air mixture to create the power expansion. Care has to be taken to avoid spontaneous ignition of the natural gas–air mixture, but with careful design the engine can operate at close to the compression conditions of a diesel engine, with a high-power output and high efficiency, yet with the emissions close to those of a gas-fired spark-ignition engine. However, efficiency tends to fall and emissions of unburned hydrocarbons and carbon monoxide rise at part load.

Typical dual-fuel engines operate with between 1% and 15% diesel fuel. Since a dual-fuel engine must be equipped with diesel injectors, exactly as if it were a diesel engine, a dual-fuel engine can also burn 100% diesel if necessary, though with the penalty of much higher emissions.

STIRLING ENGINES

Whereas fuel combustion takes place within the cylinders of an internal combustion engine, the heat energy used to drive a Stirling engine is applied outside the cylinders, which are completely sealed. The engine was designed by a Scottish Presbyterian minister, Robert Stirling, who received his first patent in 1816.

The original Stirling engines used air within the cylinders and were called air engines but modern Stirling engines usually employ helium or hydrogen. Both gases are capable of absorbing a large quantity of heat rapidly, a key advantage for an external combustion engine. The gas within the closed cycle system will be pressurized to up to 20 MPa, roughly 200 times atmospheric pressure.

The simplest form of Stirling engine consists of two cylinders, each sealed with a piston but with the sealed end of each cylinder linked so that the working fluid, which is sealed within the two cylinders, can move from one to the other during the cycle (Figure 5.5). The heat source that provides the energy to drive the cycle is applied to one cylinder while a cold sink, usually ambient air, is applied to the second. The two cylinders are linked, externally, too, by a system of levers that ensures their movements are synchronized. There will usually be a

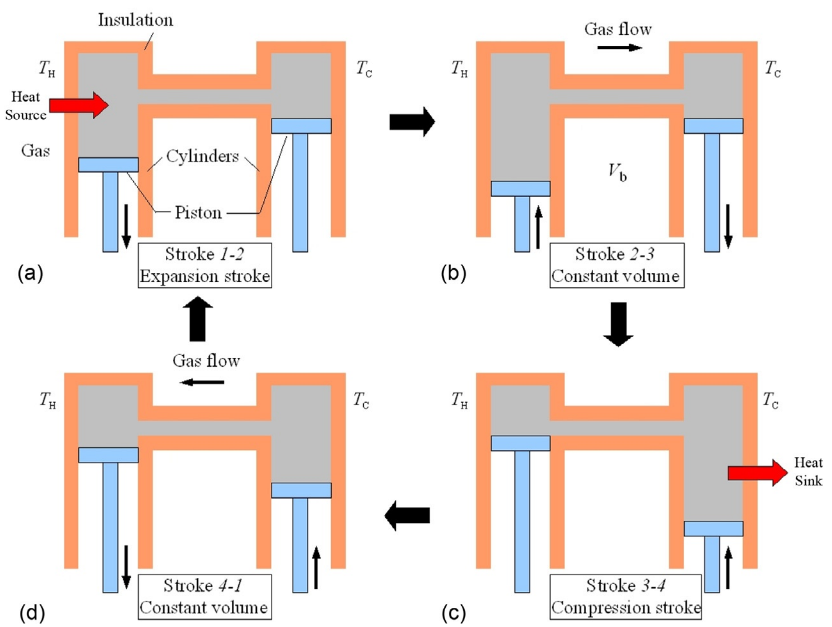

FIGURE 5.5 Stirling engine cycle.

flywheel attached to the system to store energy from the power stroke, one of four in the cycle, in the form of angular momentum to drive all four stages.

In the first part of the cycle, the power stroke for the working fluid is contained in the hot cylinder and the cold cylinder piston is at the top of its cylinder, which is empty. The heat that is continuously applied to the hot cylinder causes the working fluid to heat up and increase in pressure and expand, forcing the piston to move toward the bottom of the cylinder. During this part of the cycle the cold cylinder piston remains stationary. The expansion of the hot gas provides the power that drives the engine through its whole cycle.

In the second stage, the cold cylinder piston starts to move toward the bottom of its cylinder while at the same time the hot cylinder piston moves to the top of its cylinder, causing the working gas to move, essentially at constant pressure, from one cylinder to the other. However, as the gas is drawn into the cold cylinder it becomes cooled and the pressure falls.

In the third stage of the cycle the cold cylinder piston starts to move toward the top of its cylinder again, compressing the cold gas and generating some heat in the process, which is removed by the cold sink. During this stage the hot cylinder piston remains stationary.

The final stage, another essentially constant pressure stage, involves the cold cylinder moving to the top of its cylinder, forcing all the working fluid into the hot cylinder, while the piston in the hot cylinder starts to move down to accommodate the cold working fluid. As this begins to become heated by the hot source, the cycle starts again. In essence, the full cycle involves the expansion and then the contraction of the working fluid.

Actual Stirling engines are generally more complex than this. Some replace one cylinder by a device called a displacer, and the working fluid often passes through a recuperator between the hot and cold cylinders so that energy otherwise lost can be captured and reused. This allows the cycle to approach the theoretical Carnot cycle efficiency for a heat engine, and high efficiency is one of the great attractions of the Stirling engine.

The other great advantage of the Stirling engine is that the heat energy is applied externally. Thus the energy can, in theory, be derived from any heat source. Stirling engines have been used to exploit solar energy and for biomass applications. However, their use is not widespread. Typical engine sizes in use and development range from 1 kW to 150 kW. In solar thermal applications a Stirling engine could theoretically achieve close to 40% energy conversion efficiency. The best so far recorded is just over 31%, which is still high for solar conversion.

COGENERATION

When an internal combustion engine is used to generate electricity, a large part of the energy supplied to the engine in the form of fuel emerges as heat in the exhaust from the engine or is dumped into the atmosphere by engine cooling systems. If this heat can be captured it can be utilized for space heating or for heating water, potentially making the energy usage much more efficient.

The efficiency of piston engine–based power generation varies from 25% for small engines to close to 50% for the very largest engines. This means that between 50% and 75% of the fuel energy actually emerges as waste heat. There are four primary sources of waste heat in an internal combustion engine: engine exhaust, engine case cooling water, lubrication oil cooling water, and, where one is fitted, turbocharger cooling.[4] Each of these can be used as a source of heat in a reciprocating engine cogeneration system.

The exhaust gas contains up to one-third of the fuel energy and 30–50% of the total waste heat from the engine. Exhaust heat is not normally captured but it is straightforward to fit a heat-recovery system to the exhaust of an engine if the heat is required. The exhaust temperature is typically between 370 °C and 540 °C. This is sufficiently high that it can be used to generate medium-pressure steam if required. Otherwise, it can be used to generate hot water.

The main engine case cooling system can capture up to 30% of the total energy input. Cooling water exits the cooling system at up to 95 °C.[5] In a cogeneration system this will be passed through a heat exchanger to provide a source of hot water. Engine oil and turbocharger cooling systems will provide additional energy that can also be used to supply hot water.

If all the heat from the exhaust and cooling systems of an engine is exploited, around 70–80% of the fuel energy can be used. However, this can generally only be fully exploited when there is a need for hot water. Engine exhaust gases have also been used directly for drying in some situations.

Since cooling systems are fitted to internal combustion engines whether the waste heat is exploited or not, the use of these systems in combined heat and power applications offers a logical extension of their application. Cogeneration systems based on small engines can provide power, space heating, and hot water to homes and commercial offices, while large engines can produce power and low-grade process heat for small industrial operations. The economics of these systems can be quite favorable where there is a use for the waste heat. As a consequence, the cogeneration market, particularly for small systems, is buoyant and is likely to become more so if fuel costs continue to rise.

COMBINED CYCLE

The waste heat from the exhaust of an internal combustion engine is generally hot enough to generate medium-pressure steam. In the case of small engine installations, steam production is not normally an economical option unless

4. A turbocharger is sometimes used to compress air before it is admitted into the cylinder of an internal combustion engine. This can lead to improved performance by generating greater power from the engine.

5. It may be hotter if the cooling system is pressurized.

there is a local use for low-quality steam. In the case of a large diesel installation, however, the engine exhaust can be used to generate steam in a boiler, steam that can drive a steam turbine to produce additional energy. This forms the core of a diesel engine–based combined cycle plant.

Diesel engine combined heat and power systems are rare because they are generally only economical on very large engines. Typical of this sort of application is a generating plant that was installed in Macau in 1987. This plant was equipped with a slow-speed diesel engine with a capacity of 24.4 MW. The engine exhaust was fitted with a waste heat boiler and steam turbine that could generate an additional 1.34 MW when the engine was operating at full power, thus contributing around 5% of the plant output. As a result of this and other measures a fuel-to-electricity conversion efficiency of close to 50% was achieved.

Large engines of this type are frequently derived from marine engines and the original engines upon which they are based are not normally optimized for combined cycle operation. In particular, the cooling system is designed to keep the engine as cool as possible. For best combined cycle performance, however, it is preferable to run the engine as hot as possible, because the higher the exhaust gas temperature, the more efficient the steam turbine cycle. High-temperature operation can also improve engine efficiency because the potential thermodynamic efficiency will increase with operating temperature.

Combined cycle performance of a large diesel engine can, therefore, be improved by modifying engine components such that they can operate continuously at a higher temperature. Such modifications may require more expensive materials capable of withstanding the more extreme conditions. For example, the top of the piston may be made from an alloy that allows it to remain uncooled while exhaust valves are treated with advanced coatings able to resist the high exhaust gas temperature.

These modifications allow a higher temperature exhaust that can be used to generate higher-quality steam to drive a steam turbine. With these measures it may be possible to achieve a fuel-to-electricity conversion efficiency of close to 55%. This is the efficiency target for a plant in Wasa, Finland, installed in 1998. The plant has two 17 MW diesel engines and a single steam turbine. Efficiency in this case is improved by using seawater cooling for the steam turbine condenser. The additional expense of the waste heat recovery and steam turbine will generally only prove cost effective if the engine is to be used for base-load operation.

EMISSION CONTROL

Piston engine power units generally burn fossil fuels, and the environmental considerations that need to be taken into account are exactly the same as those that affect all coal-, oil-, and gas-fired power plants—that is, all the emissions resulting from fuel combustion. In the case of internal combustion engines the main emissions are nitrogen oxide, carbon monoxide, and volatile organic compounds (VOCs). Larger diesel engines, particularly those burning heavy diesel fuel, will also produce particulate matter and some sulfur dioxide.

Nitrogen oxide is formed primarily during combustion by a reaction between nitrogen and oxygen in the air mixed with the fuel. This reaction takes place more rapidly at higher temperatures. In lean-burn gas engines where the fuel is burned with an excess of air, temperatures can be kept low enough to maintain low nitrogen oxide emissions. The diesel cycle depends on relatively high temperatures, and as a consequence of this, produces relatively high levels of nitrogen oxide. Table 5.3 compares emissions from the two types of engine.

When the fuel in an internal combustion engine is not completely burned the exhaust will contain both carbon monoxide and some unburned hydrocarbons. Carbon monoxide is hazardous at low levels and its emissions are regulated like those of nitrogen oxide. Unburned hydrocarbons are classified as VOCs and their emissions are also controlled by legislation.

Natural gas contains negligible quantities of sulfur so gas engines produce no sulfur dioxide. Diesel fuels can contain sulfur. Small- and medium-size diesel engines generally burn lighter diesel fuels that contain little sulfur. Larger engines can burn heavy residual oils that are comparatively cheap but often contain significant levels of sulfur. Since sulfur can damage the engine, it is normal to treat this type of fuel first to remove most of the sulfur.

Liquid fuels may produce particulate matter in an engine exhaust, the particles derived from ash and metallic additives. Incomplete combustion of heavy fuel can also lead to the emissions of particulate matter.

Nitrogen Oxide

The most serious exhaust emissions from a piston engine is nitrogen oxide. Engine modifications that reduce the combustion temperature of the fuel, such as the use of a lean-fuel mixture described earlier, provide the first step in reducing these emissions. Natural gas engines designed to burn a very lean fuel (excess air) provide the best performance: 45–150 ppmV or 1–3 g/kWh. Diesel engines present a greater problem because of the higher combustion temperature (Table 5.4).

An additional technique that is being applied to internal combustion engines to reduce nitrogen oxide is exhaust gas recirculation. This involves taking some of the exhaust gas and mixing it with the air used to feed the engine. The effect is to reduce the overall oxygen concentration and thereby reduce the combustion temperature. This reduces nitrogen oxide production but will also reduce engine efficiency. However, depending on the regulatory regime, all types of internal combustion engines may require some additional form of post-combustion nitrogen oxide emission control.

For small gasoline engines a simple catalytic converter of the type used in automobiles is often the most effective solution. However, this type of system cannot be used with diesel or with lean-burn engines. New catalysts for use with lean-burn engines are currently under development. Where a catalytic converter can be used, nitrogen oxide reduction is around 90% or more.

Automobile-style catalytic converters are a relatively expensive means of reducing nitrogen oxide emissions. For large engines, the more economical alternative is to use a selective catalytic reduction (SCR) system that can be applied to both stationary engines and large truck engines. An SCR system also employs a catalyst, but in conjunction with a chemical reagent, normally ammonia or urea, which is added to the exhaust gas stream before the emission control system. The reagent and nitrogen oxide react on the catalyst and the nitrogen oxide is reduced to nitrogen. This type of system will reduce emissions by 80% to 90%. However, care has to be taken to balance the quantity of reagent added so that none emerges from the final exhaust to create a secondary emission problem.

Carbon Monoxide, VOCs, and Particulates

The emission of carbon monoxide, VOCs, and some particulate matter can be partially controlled by ensuring that the fuel is completely burned within the engine. This is simplest in lean-burn engines but conditions within these engines compromises efficiency. With all engines, careful control of engine conditions and electronic monitoring systems can help maintain engine conditions at their optimum level. Old engines as they become worn can burn lubrication oil, causing further particulate emissions.

For larger engines, particularly diesel engines, engine control systems also will not maintain emissions sufficiently low to meet statutory emission standards. In this case an oxidation catalyst will be needed to treat the exhaust gases. When the hot gases pass over the oxidation catalyst, carbon monoxide, unburned hydrocarbons, and carbon particles are oxidized by reacting with oxygen remaining in the exhaust gases, completing the combustion process and converting all the materials into carbon dioxide.

Sulfur Dioxide

Sulfur emissions can be found in diesel engines that burn fuel containing sulfur. Many engines now burn low-sulfur fuels with less than 0.05% sulfur content. However, some diesels and the heavy fuel oils that very large engines burn may contain significant amounts of sulfur. The latter may contain as much as 3.5% sulfur. The best way of controlling sulfur emissions from internal combustion engines is to remove the sulfur from the fuel before use. However, in the worst case a sulfur capture system can be fitted. This is likely to be similar to the scrubbing tower used in a coal-fired power plant but at a much smaller scale. The use of such a system adds to both capital and maintenance costs and affects plant economics. It is only likely to be cost effective in the very largest reciprocating engine-based power plants.

Carbon Dioxide

Internal combustion engines, in common with all heat engines that burn carbon-based fuel, generate carbon dioxide, which is released in the exhaust gases leaving the engine. The relative amount produced during electricity generation depends on the efficiency of the engine. A large, high-efficiency diesel engine operating at close to 50% efficiency will produce significantly less carbon dioxide for each unit of electricity it generates than a small gasoline engine operating at perhaps 20% efficiency.

Currently the only way of effectively eliminating carbon dioxide emissions from such engines is to run them on a biofuel such as ethanol or biodiesel that has been derived from plants. The principle here is that although the combustion of the fuel will still produce carbon dioxide, the regrowth of the plants that were used to produce the fuel will absorb the same amount of carbon dioxide from the atmosphere, so that for a full cycle of growth, fuel production and combustion, the net amount of carbon dioxide added to the atmosphere is zero.

Research is underway to develop systems to capture carbon dioxide from the exhaust of reciprocating engines and a variety of techniques are being explored based on some form of post-combustion capture. Whether such systems will ever be used extensively on reciprocating engines seems doubtful since the cost is likely to be prohibitive. However, for very large power generating systems it might eventually be both technically feasible and economical.

COST OF RECIPROCATING ENGINE-BASED POWER GENERATION

The economics of power generation based on reciprocating engines depends to a large extent on the use to which the engine is to be put. The cheapest engines available are small petrol-driven engines based on car engines, which are manufactured in large numbers each year. These engines can be purchased as standby generators for as little as $250/kW. Such engines are cheap so they are well suited to applications where they will only be required to operate infrequently. However, they are expensive to run since their energy conversion efficiency is relatively low and they have short lifetimes, requiring the extensive and regular maintenance of an automotive engine.

Large engines designed for power generation are generally much more expensive. Natural gas–fired engines of around 300 kW are likely to cost around $2000/kW. While detailed cost data tends to be proprietary, evidence suggests that capital costs drop as engine size rises into the megawatt and multi-megawatt range. All these larger engines are built to be able to operate for long periods between maintenance. They are generally more efficient than the smaller engines too, so their operating costs are lower.

Table 5.5 shows figures for a range of generic power generation systems based on reciprocating engines that illustrates these trends. The costs in this

TABLE 5.5 Cost and Efficiency Figures for a Series of Reciprocating Engine Systems

	Capacity (KW)	Total Installed Cost ($/kW)	Engine Speed (rpm)	Efficiency (HHV, %)
System 1	100	2210	1800	28.4
System 2	300	1940	1800	34.6
System 3	800	1640	1800	35
System 4	3000	1130	900	36
System 5	5000	1130	720	39

Source: *Technology Characterization: Reciprocating Engines,* U.S. Environmental Protection Agency, 2008.

table are in year 2010 U.S. dollars. The smallest system in the table, with a generating capacity of 100 kW, is based on a high-speed engine operating at 1800 rpm. Efficiency is 28.4% and the cost is $2210/kW. At the other end of the scale a 5 MW system is based on a medium-speed engine running at 720 rpm. This has an efficiency of 39% and an installed cost of $1130/kW.

The cost effectiveness of most systems such as those in Table 5.5 will depend on whether they can be used to supply heat as well as electrical power. All the systems in the table were assessed for their cogeneration efficiency when providing hot water. The most efficient were systems 2 and 3 with overall efficiencies of 78% and 79%, respectively. The least efficient was system 4 at 73%.

The cost of Stirling engines is much higher than for most internal combustion engines because they are not produced in large enough quantities to bring about the economy of volume production. Estimates vary widely, from as low as $2000/kW to as high as $50,000/kW.

Chapter 6

Combined Heat and Power

The production of electricity from coal, oil, gas, and biomass is an inefficient process. While some modern combustion plants can achieve 60% energy conversion efficiency, most operate closer to 30%, and smaller or older units may reach only 20%. The United States, which has a typical developed-world mix of fossil fuel–based combustion plants, achieves an average power plant efficiency of 33%, a level that has barely shifted for the past 30 years. Other countries would probably struggle to reach even this level of efficiency.

Put another way, between 40% and more than 80% of all the energy released during combustion in power plants is wasted. The wasted energy emerges as heat that is dumped in one way or another. Sometimes it ends up in cooling water that has passed through a power plant and then returned to a river or the sea, but most often it is dissipated into the atmosphere through some form of air–heat exchanger. This heat can be considered a form of pollution.

Efficiency improvements can clearly curtail a part of this loss. But even with the most efficient energy conversion system, some loss of energy is inevitable. Neither thermodynamic nor electrochemical energy conversion processes can operate even theoretically anywhere near 100% efficiency and practical conversion efficiencies are always below the theoretical limit. So while technological advances may improve conversion efficiencies, a considerable amount of energy will always be wasted.

While this energy cannot be utilized to generate electricity, it can still be employed. Low-grade heat can be used to produce hot water or for space heating,[1] while higher-grade heat will generate steam that can be exploited by some industrial processes. In this way the waste heat from power generation can replace heat or steam produced from a high-value energy source such as gas, oil, or even electricity. This represents a significant improvement in overall energy efficiency.

Systems that utilize waste heat in this way are called combined heat and power (CHP) systems (the term *cogeneration* is often used too). Such systems can operate with an energy efficiency of up to 90% when heat usage is taken into account. This represents a major savings in fuel cost and in overall

1. Heat can also be used to drive chillers and cooling systems. These are not considered separately here.

environmental degradation. Yet, while the benefits are widely recognized, the implementation of CHP remains low.

Part of the problem lies in the historical and widespread preference for large central power stations to generate electricity. Large plants are efficient and they are normally built close to the main transmission system so that power can be delivered into the network easily. They may also be sited close to a source of fuel. This will often mean that they are far from consumers that can make use of their waste heat.

If central power plants are built in or near cities and towns then they can supply heat as well as power by using their waste heat in district heating systems. Municipal utilities in some European and U.S. cities have in the past built power plants within cities they serve to exploit this market for heat and power, but it is not an approach that has been widely adopted and environmental considerations makes building large power plants in cities more difficult today. There are also many examples of power plants being built close to industrial centers such that they can provide high-grade steam for industrial use. In the main, however, large fossil fuel power plants simply waste a large part of the energy they release from the fuel.

At a smaller scale, the situation is slightly better. At the distributed generation level, in particular, where power is generated either for private use or to feed into the distribution level of a power supply network, it is much easier to find local sources of heat demand that can be met at the same time as power is generated. This means that there are greater opportunities to achieve higher energy efficiency.

In an energy-constrained and environmentally stressed world energy efficiency represents one of the best ways of cutting energy use and reducing atmospheric emissions. The German government has estimated that 50% of its electricity could be supplied through CHP systems. There are economic advantages too that make greater use of CHP an extremely attractive proposition. In spite of these arguments, growth in the use of CHP has been painfully slow and it remains a major challenge for the electricity industry to achieve higher energy efficiency through the use of CHP.

HISTORICAL BACKGROUND FOR COMBINED HEAT AND POWER USAGE

The concept of combined heat and power generation is not new. Indeed, the potential for combining the generation of electricity with the generation of heat was recognized early in the development of the electricity-generating industry. In the United States, for example, at the end of the 19th century city authorities used heat from the plants they had built to provide electricity for lighting to supply hot water and space heating for homes and offices too. These district heating schemes, as they became known, were soon being replicated in other parts of the world.

In the United Kingdom, around that time, a small number of engineers saw in this a vision of the future. Unfortunately, their vision was not shared and uptake was slow. It was not until 1911 in the United Kingdom that a district heating scheme of any significance, in Manchester city center, was developed.[2] Fuel shortages after World War I, followed by the Great Depression, made district heating more attractive as electricity generation expanded in Europe during the 1920s and 1930s. Even there, however, take up was patchy. Nevertheless, by the early 1950s district heating systems had become established in some cities in the United States, in European countries such as Germany and Russia, and in Scandinavia. In other countries like the United Kingdom there was never any great enthusiasm for CHP and it gained few converts, though a number of schemes were built after World War II as regions devastated by bombing were rebuilt.

This pattern of patchy exploitation has continued and the situation is complicated by the fact that it is almost impossible, economically, to build district heating infrastructure in modern cities that lack it. The centralization of the electricity supply industry must take some blame for this lack of implementation. Where a municipality owns its own power-generating facility it can easily make a case on economic grounds for developing a district heating system. But when power generation is controlled by a centralized, often national body, the harnessing of small power plants to district heating networks can be seen as hampering the development of an efficient national electricity system based on large, central power stations—unless, that is, the CHP approach is already a part of the philosophy of the national utility.

Power industry structure is not the only factor. Culture and climate are also significant. So, while countries such as the United Kingdom failed to make significant investment in district heating, Finland invested heavily. Over 90% of the buildings in its major cities are linked to district heating systems, and over 25% of the country's electricity is generated in district heating plants. Many Russian cities, too, have district heating systems with heat generated from large local power stations. Even some nuclear power plants in Russia are harnessed in this way.

District heating was—and remains—a natural adjunct of municipal power plant development. But by the early 1950s the idea was gaining ground that a manufacturing plant, like a city, might take advantage of CHP too. If a factory uses large quantities of both electricity and heat, then installing its own power station allows it to control the cost of electricity and to use the waste heat produced, to considerable economic benefit. Paper mills and chemical factories are typical instances where the economics of such schemes are favorable and many such plants operate their own CHP plants.[3]

2. Stewart Russell, Combined Heat and Power in Britain, in *The Combined Generation of Heat and Power in Great Britain and the Netherlands: Histories of Success and Failure*, R1994: 29 (Stockholm: NUTEK, 1994).

3. Applications of this type are frequently designated cogeneration rather than combined heat and power. The underlying premise is identical, however.

While this idea slowly gained ground, technological advances during the 1980s and 1990s made it possible for smaller factories, offices, and even housing developments to install CHP systems. In many cases this was aided by the deregulation of the power supply industry and the introduction of legislation that allowed small producers to sell surplus power to the local grid. Since the middle of the 1990s the concept of distributed generation has become popular and this has also encouraged CHP.

Recent concern for the environment now plays its part too. Pushing energy efficiency from 30% to 70% or 80% more than halves the atmospheric emissions from a power station on a per-kWh basis. Thus, CHP is seen as a key emission control strategy for the 21st century. But while environmentalists call for expanded use, actual growth remains slow.

COMBINED HEAT AND POWER PRINCIPLES AND APPLICATIONS

From single home heat-and-electricity units to municipal power stations supplying heat and power to a city, from paper mills burning their waste to provide steam and heat to large chemical plants installing gas turbine–based CHP facilities, CHP installations can be as different as their applications are varied. However, they do share one theme. Ideally, the heat and electricity from a CHP plant will be supplied to the same users. While this is not an absolute requirement it is a pragmatic principle for a successful CHP scheme.

If the heat and power plant is supplying both types of energy to the same users, be they an industrial plant or households, then the economics of the plant will remain sound so long as the customers remain. However, if a plant supplies electricity to one customer and heat to another, it can be undermined by the loss of either. Part of this risk is mitigated if a plant can export electricity directly into the grid, but, in general, the economics of CHP will be most soundly based where the same customers take both.

Examples of how this can be achieved exist at all levels within the electricity market from the smallest electricity consumer to the largest. At the very bottom of the scale many home owners buy electricity from a utility and either use this for heating or to purchase natural gas to provide hot water for domestic use and space heating. However, it is now possible to install domestic CHP systems based on micro-turbines or fuel cells that will replace a domestic hot water system and at the same time generate electricity for use by the household, with excess power perhaps being sold to the local utility. While such systems remain costly, they are likely to become more cost effective in time.

On a larger scale a micro-turbine or a reciprocating engine burning natural gas can be used to supply both electricity and heat to an office building, a large block of apartments, or a small commercial or industrial enterprise. Such systems are widely available and can be installed in urban environments with ease. The system may be connected in such a way that excess power can be exported

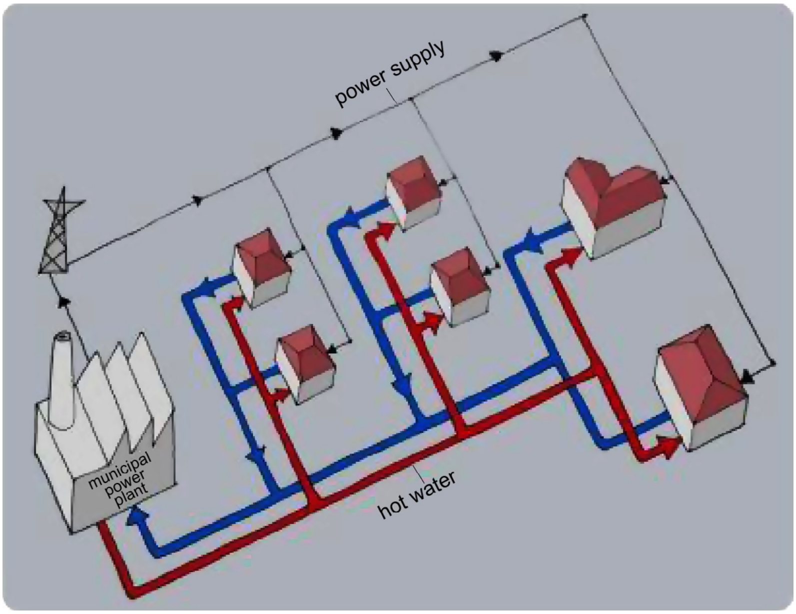

FIGURE 6.1 Plan of a typical municipal combined heat and power system.

to the grid, although the sizing of such systems is normally based on heat demand; power demand will be generally higher than the system can supply alone.

At the top end of the capacity scale, a municipal power plant based on a coal-fired boiler or a gas turbine can provide electricity for a city and heat for that city's district heating system (Figure 6.1). As already noted, new systems of this type are difficult to install in existing cities but opportunities arise where major redevelopment takes place and new urban housing schemes can be built with district heating too.

In all these cases the supply of electricity is the primary requirement, but the use of heat from the power generation facility improves the economics considerably while offering significant environmental benefits. Similar opportunities exist in industry, but in many industrial cases the situation is reversed and it is heat or steam that is the primary requirement with electricity production a secondary consideration. However, the same arguments apply. There are many industrial processes that require a source of heat and all industrial plants will use some electricity. Often the two can be combined to good economic effect once the benefits are recognized. So where, in the past, a paper manufacturer would have installed a boiler to supply heat while buying power from the grid, now the same manufacturer is more likely to install a CHP plant, often fired with waste produced by the paper manufacturing process. Chemical plants often require a supply of high-grade heat too, as do some refineries. Food-processing plants may require large heat supplies as well.

Such instances represent the ideal but a good match of heat and power demand is not always possible at this level. However, provided the scale is large enough, such plants may be able to install a grid-connected power plant that exports electricity while using heat generated from the same plant to supply its industrial needs. Alternatively, a large CHP plant can be built to supply an industrial site supporting a range of different industries, some with large heat requirements, others with demand for electrical power.

Domestic heat consumption remains the biggest challenge. Where district heating networks exist, a good balance between domestic heat and electricity demand is possible. But where these do not exist the only solution is either power stations meeting the electricity demand only, or domestic CHP systems. The latter are the only solution for many households but their installed base is still very small and, as already noted, cost is high. If costs can be brought down then such systems offer a real chance of a significant change in global energy usage.

CHP TECHNOLOGY

All types of power generation technology that are based on heat engines are capable of being integrated into a combined heat and power system. Thus, all fossil fuel–fired power plants and biomass power plants can be adapted to create CHP systems. In addition, electrochemical fuel cells can be an excellent source of heat and power. Among renewable technologies, solar thermal power plants can provide heat in addition to electricity if necessary, and geothermal energy has been widely exploited for both electricity generation and district heating. Other renewable technologies like wind, hydropower, and marine power involve no heat generation. However, solar photovoltaic power generation can, in principle at least, be exploited in conjunction with solar heating because the solar cell only uses a part of the incident light and rejects most of the heat-bearing infrared radiation. Nuclear power, too, can be exploited for heat and power generation although its use is rare.

While this range of technologies offers a wide choice for a CHP plant, the type of heat required from a CHP application will often narrow the choice of technology. If high-quality steam is demanded then a source of high-temperature waste heat will be needed. This can be taken from a steam turbine–based power plant, it can be generated using the exhaust of a gas turbine, and it can be found in a high-temperature fuel cell. Other generating systems such as piston engines or low-temperature fuel cells are only capable of generating hot water, and perhaps low-quality steam.

The way in which a CHP plant is to operate is another important consideration. Which of the two types of energy—heat or electricity—will take priority? If heat is the most important consideration, particularly if this is high-grade industrial heat, then a system based on a steam-generating boiler and steam turbine will probably be most appropriate. The boiler will be sized to meet the maximum steam demand while a steam turbine is available to exploit excess steam to generate electricity. If there is sizable demand for both heat and electricity, then either a steam turbine of a gas turbine–based system might be the most suitable, with exhaust heat from the gas turbine used to generate steam and a steam turbine to exploit any excess. However, this will require an electricity demand at least equal to the output of the gas turbine generator.

For smaller applications and where only hot water is needed, a reciprocating engine, micro-turbine, or low-temperature fuel cell might offer the best match. Again, however, the mode of operation will determine the optimum choice. If the unit is to supply power to a particular consumer or group of consumers, with its output following their demand, then a generating unit that can operate efficiently at different load levels such as a piston engine or fuel cell will probably be the best solution. However, if it is going to provide base-load generation then part-load efficiency will be of less significance.

Finally, location will be important. It will not be possible to install some types of CHP plant in urban areas because of the emissions and the noise they generate. Therefore, this will limit the technologies available for use in this situation.

The quantity of heat that will be available will vary from technology to technology. Table 6.1 gives typical energy conversion efficiency ranges for modern fossil fuel–burning power plants. However, not all the energy that is not converted into electricity will be available as heat. Modern high-efficiency coal-fired power plants can operate between 38% and 45% efficiency, though there are many that are much less efficient. However, high-efficiency coal-fired plants produce little usable waste heat unless overall efficiency is compromised, because the steam exiting the steam turbine is generally at a very low temperature and pressure. Gas turbines provide more flexibility while offering a similar electrical energy conversion efficiency because their exhaust gases can provide high-grade heat.

TABLE 6.1 Typical Power Plant Energy Conversion Efficiencies

Type of Plant	Efficiency
Conventional coal-fired power plant	38–45% for modern high-efficiency coal plants
Heavy industrial gas turbine	Up to 42%
Aero-derivative gas turbine	Up to 46%
Fuel cell	25–60% depending on type
Natural gas–fired reciprocating engine	28–42%
Diesel engine	30–50%

Where even more flexibility is required, it is possible to design a plant to produce less electricity and more heat than the efficiency figures in Table 6.1 suggest. Some technologies are amenable to this strategy, others are not. Most flexible are boiler/steam turbine plants, but gas turbine CHP units can easily be adapted in this way too.

All the technologies employed in CHP plants have their own chapters in this book where detailed accounts of their operation can be found. In discussing these technologies here, consideration will only be given to factors of specific relevance to CHP. Please refer to the other chapters for fuller accounts of each technology.

PISTON ENGINES

There are two primary types of piston engine for power generation: the diesel engine and the spark-ignition gas engine. Of these the diesel engine is the most efficient, reaching close to 50% energy conversion efficiency. The spark-ignition engine burning natural gas can achieve perhaps 42% efficiency but it is much cleaner than the diesel. The level of emissions from an uncontrolled diesel engine are such that it is impossible to obtain authorization to use a diesel engine for continuous power generation service in some parts of the world unless it is fitted with an extensive emission control system. Natural gas engines can often operate with minimum emission control.

Piston engines are well suited to CHP applications where hot water is required because much of the energy that is not converted into electricity appears as heat at a temperature suitable for water heating. There are four sources of heat in a piston engine: the engine exhaust, engine jacket-cooling system, oil-cooling system, and turbocharger cooling system (if a turbocharger is fitted), as shown in Figure 6.2. Engine exhaust can provide low- to medium-pressure steam and the engine jacket-cooling system can provide low-pressure steam. However, for most piston engine CHP applications all the sources of heat

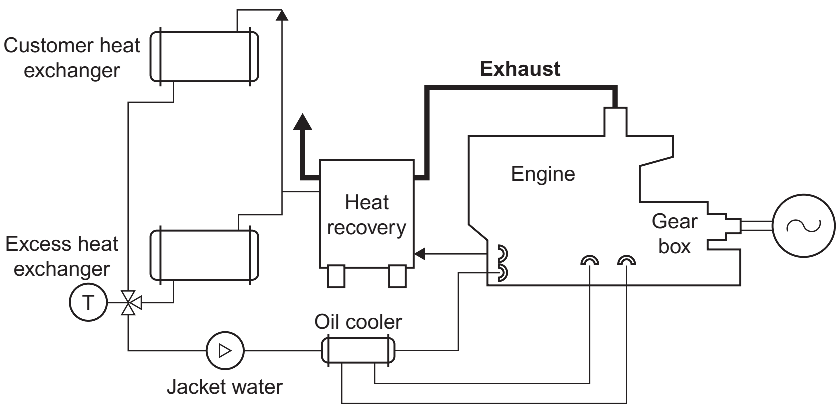

FIGURE 6.2 Piston engine CHP.

are used to generate hot water rather than steam. If all four sources of heat are exploited, roughly 70–80% of the energy in the fuel can be utilized.

Piston engine power plants suitable for general CHP applications are available in sizes ranging from a few kilowatts to 6.5 MW. These engines are particularly good at load following; a spark-ignition engine efficiency falls around 10% at half load while diesel engine efficiency barely drops over this range. There is no significant penalty in terms of engine wear for variable-load operation either, unlike for some other engine types.

Applications for piston engine CHP plants include small offices and apartment blocks, hospitals, government installations, colleges, and small district heating systems. Engines tend to be noisy, so some form of noise insulation is normally required. Emissions of gas engines can normally be controlled with simple exhaust catalytic converter systems, but diesel engines usually require much more elaborate measures to control their higher nitrogen oxide and particulate emissions.

STEAM TURBINES

A steam turbine is one of the most reliable units for power generation available. Large utility steam plants designed exclusively for power generation have efficiencies up to 45%, but smaller units employed for CHP applications generally provide efficiencies of 30–42%. These turbines are usually simpler in design too. Steam turbines are available in virtually any size from 50 kW to 1300 MW.

A steam turbine cannot generate electricity without a source of steam. This is normally provided by a boiler in which a fossil fuel or biomass fuel is burned. This makes steam turbine CHP extremely flexible because the power and steam generation are essentially independent of one another. A steam turbine will normally be used in a CHP system only where there is a demand for high-quality, high-pressure steam for some industrial process.

There are a variety of ways of configuring a boiler/steam turbine system to provide both electrical power and heat. One method is to take heat directly from the boiler to supply heat to whatever process needs it with any surplus being directed to a steam turbine to generate electricity. Such an arrangement normally will be economically effective if most of the steam is being used by the industrial process; the addition of a small steam turbine unit then allows limited power generation when excess steam is available.

A more common configuration utilizes what is known as a back-pressure steam turbine (Figure 6.3). In this configuration the steam from the boiler goes directly to the steam turbine, which extracts a part of the energy contained within it. The steam exiting the turbine, still at an elevated temperature and pressure, is then directed to the process where heat is required. This will normally be used when lower-pressure and lower-temperature steam is required, but by balancing the size of the turbine and boiler, steam temperature and pressure can be

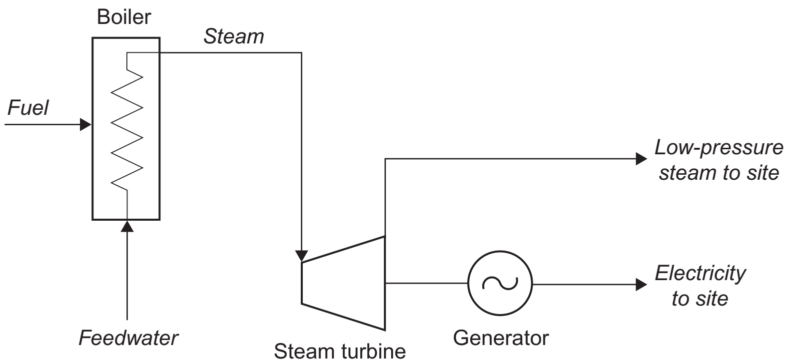

FIGURE 6.3 Steam turbine CHP.

tailored to suit the process in question. Back-pressure steam turbine CHP systems are widely used in industry.

A third approach is to feed the steam from the boiler into a condensing turbine (one that condenses steam to water at its exhaust to create the largest temperature and pressure drop possible), but then extract steam from a point partway through the steam turbine, where steam temperature and pressure match the heat requirement. Turbines designed for operation in this way are called extraction steam turbines and are commonly found in industrial plants.

These different configurations allow for considerable flexibility when designing a steam turbine–based CHP system. Systems can be more complex if an industrial facility has several sources of steam. These may feed surplus steam into different stages of a single steam turbine to extract the most energy possible.

Steam turbine CHP systems will generally be relatively large, from a few megawatts to perhaps hundreds of megawatts. Most often such plants will be associated with a single industrial plant or process. Depending on the fuel burned in the plant boiler to generate steam, emission control systems will be needed to limit atmospheric pollutants. Steam turbine CHP plants may burn coal or biomass but they will often burn natural gas too, which requires less by way of emission control.

GAS TURBINES

Unlike steam turbines, gas turbines burn fuel directly. Large industrial gas turbines operate with energy conversion efficiencies up to 42%, while smaller aero-derivative gas turbines can operate with an efficiency up to 46%. Gas turbine–generating capacities typically range from 1 MW up to nearly 400 MW. Units of any size can be used in CHP systems and gas turbines are probably the cheapest prime movers available today. The output of a gas turbine CHP plant can be modulated to suit the demand for heat and electricity, but they are probably best suited to continuous operation at or near full output if they are to achieve the best efficiency.

Gas turbines can burn a variety of fuels, including natural gas, distillate fuels, and biogas. However, the most popular fuel is natural gas. A stationary gas turbine is designed to generate electrical power, so there must be a market for the power generated by the plant if it is to be economically viable as part of a CHP plant. The heat output from the engine is all found it its exhaust gases, and it is these that must be exploited to provide steam or hot water. This will normally be achieved by fitting the unit with a heat-recovery steam generator. The energy contained in the exhaust of a gas turbine is suitable for generating high-pressure steam suitable for many industrial applications.

There are three main configurations for a gas turbine CHP plant. The simplest is to install a gas turbine to generate electrical power and use the exhaust gases to generate steam, all of which is used for process heat or to meet other heat demands (Figure 6.4). This basic configuration can be made more flexible in one of two ways. The first is to fit the plant with a steam turbine too, so that any excess steam can be used to generate further electrical power as in a combined cycle plant. The second is to fit the boiler with a supplementary firing system so that it can generate additional steam beyond the quantity that the gas turbine can provide.[4] There can also be a combination of supplementary firing and a steam turbine to provide more flexibility and the steam turbine

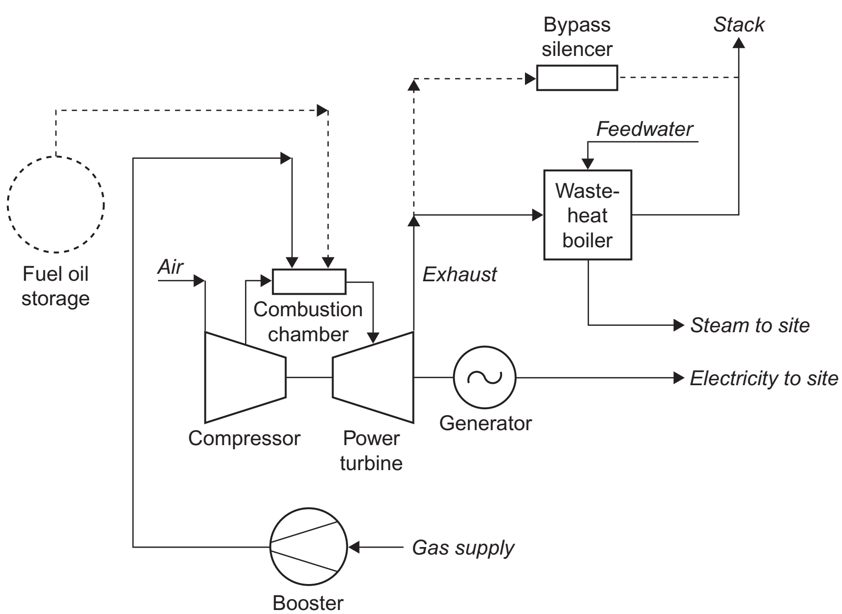

FIGURE 6.4 Gas turbine CHP.

4. The exhaust from a gas turbine contains sufficient oxygen to make it possible to install a supplementary firing system without the need for additional air.

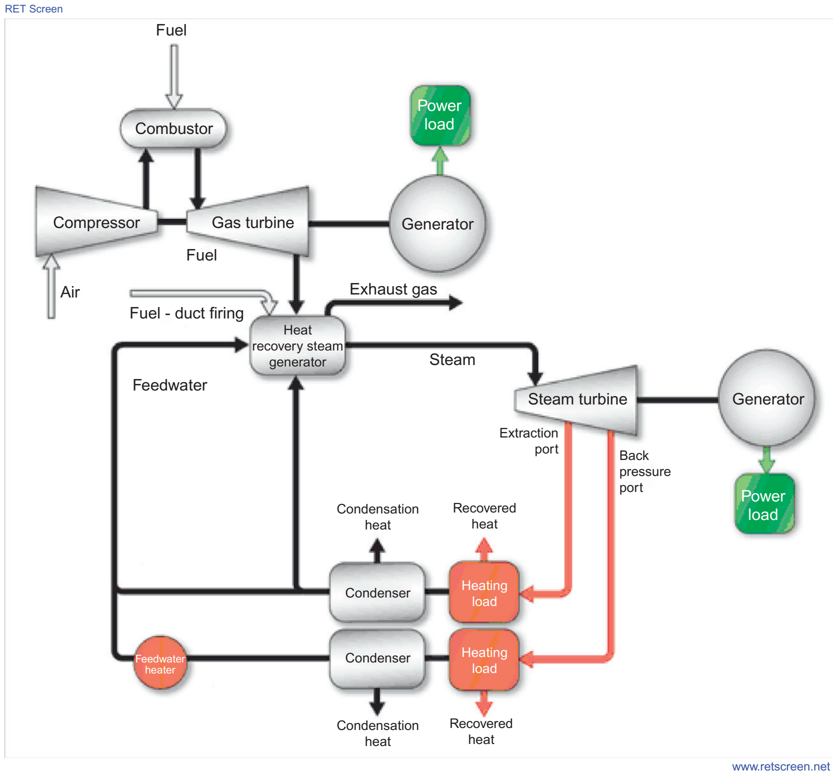

FIGURE 6.5 A complex gas turbine CHP system with an extraction steam turbine and supplementary firing.

may be a condensing steam turbine, back-pressure steam turbine, or extraction steam turbine (Figure 6.5).

Which of these configurations is chosen will depend on the balance of demand for heat and electrical power. If heat demand is relatively constant, then the simplest configuration can be chosen with the size set to produce all the heat required. Electrical output can then be supplemented if necessary from the grid. If electrical demand is high, or if power is to be exported to the grid, then one of the other configurations might make more economic sense. If the system is designed to closely match electrical and heat demand, efficiencies up to 90% can be achieved.

Gas turbine CHP systems will generally be installed in industrial situations where heat demand is high. The principal emission that needs to be controlled is nitrogen oxide. It may be possible to meet regulatory requirements without additional emission control systems when the gas turbine is small, but larger units will generally require some exhaust gas treatment system. Very large units may also need systems to control the emissions of carbon monoxide and hydrocarbons.

MICRO-TURBINES

There is another type of gas turbine called a micro-turbine that can be utilized for small cogeneration applications. Micro-turbines are tiny gas turbines that operate at very high speeds. Capacities for these machines vary from as small as 25 kW to perhaps 250 kW. There are larger units, between 250 kW and 500 kW, that are sometimes called mini-turbines, but they are best considered with traditional gas turbines and are not included in the following discussion.

Operating speed of micro-turbines can be as high as 120,000 rpm. The bearings are often air lubricated to reduce wear and most micro-turbines incorporate a generator on the same shaft to make the package as compact as possible. Units can burn gasoline, diesel, and alcohol, but for most applications they will burn natural gas.

The high operating speed means that a micro-turbine generator cannot interface directly with the grid, and most units are equipped with solid-state interfaces that convert the high-frequency output to grid frequency at 50 Hz or 60 Hz. Efficiency is relatively low for electricity generation at around 15% to 30%. This low efficiency is not generally a problem because micro-turbines are generally designed for CHP applications with waste heat-recovery systems capable of providing hot water or, in some cases, low-pressure steam.

Micro-turbines are generally supplied packaged so that all they require is a gas supply, an electrical connection, and a connection for their hot water supply. In this form they can be used in a variety of situations such as small commercial environments or office blocks. The units have low emissions so they can be deployed in urban settings without any problem. Noise generation is low too. Larger micro-turbines have been installed in schools and hospitals and they can be used in some small industrial situations.

The most recent development of micro-turbines is for domestic use. Units are packaged with an electrical output as low as 3 kW, suitable for many single homes where they will also supply hot water for heating and other uses. These domestic units are still relatively expensive but large-scale mass production could realize significant economies of scale.

FUEL CELLS

Fuel cells are electrochemical devices similar to batteries that convert the chemical energy in a fuel directly into electricity (Figure 6.6). Since fuel cells are not limited by the Carnot cycle efficiency, they are potentially more efficient than heat engines. However, very high efficiencies are difficult to realize in practice. Proton exchange membrane (PEM) fuel cells, being developed extensively for automotive applications, can achieve between 25% and 35% efficiency when operating on natural gas, although in principle these cells can run much more efficiently if they are supplied with hydrogen fuel. PEM cells operate at a relatively low temperature. High-temperature fuel cells, such as molten carbonate

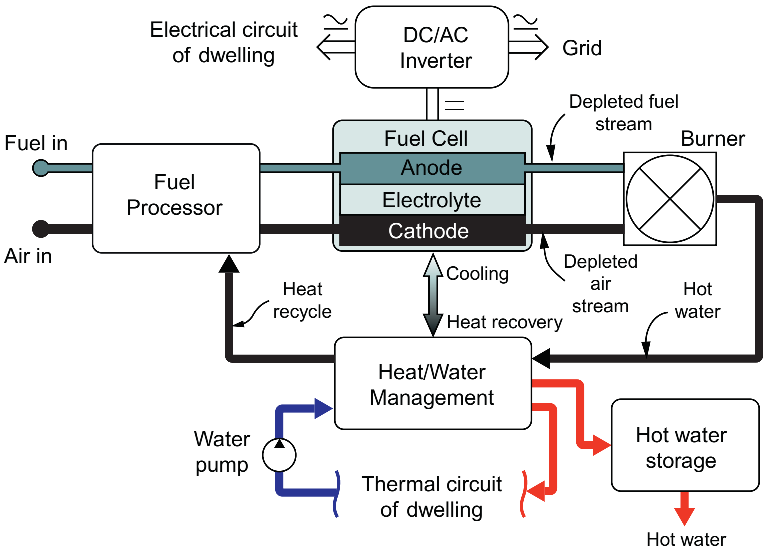

FIGURE 6.6 Fuel cell domestic CHP system.

and solid oxide fuel cells, can achieve over 40% efficiency and potentially up to 50% or higher. They are also capable of supplying high-grade heat. However, the most popular fuel cell for CHP applications is the phosphoric acid fuel cell (PAFC). Fuel cells are particularly good at load following as efficiency varies little with output.

The PAFC was one of the first fuel cells to be developed commercially. It operates at a moderate temperature of 150–200 °C, allowing it to produce low-grade heat for hot water and space heating. In common with most fuel cells, the PAFC requires hydrogen and oxygen to be supplied to its electrodes. Oxygen is provided from air, but in most applications hydrogen must be generated from natural gas in a process known as reforming. This is relatively energy intensive and reduces overall efficiency. Even so, the units can generate electricity at an efficiency around 42%.

PAFC fuel cells are packaged so they can be installed rapidly and easily, requiring only a gas supply and connections for electricity and hot water output. Typical electrical generating capacities are 100–400 kW. In CHP applications they can achieve up to 87% efficiency. Units are virtually noiseless and emissions are generally negligible making them easy to site. The range of electrical outputs available from single units makes them most suited to small commercial and institutional applications, such as schools, hospitals, and offices.

PEM fuel cells operate at a much lower temperature than PAFC fuel cells, usually around 80 °C. This means they can provide some hot water, and small units are being built aimed at the residential market. Electrical efficiency for such systems is around 30% and CHP efficiency 80%. Unit sizes can range from

3 kW to 250 kW. As with PAFC fuel cells, these cells require hydrogen to be supplied to their electrodes. Again this is normally produced by reforming natural gas.

Two main types of high-temperature fuel cells are under development: the molten carbonate fuel cell (MCFC) and the solid oxide fuel cell (SOFC). MCFC cells operate at around 650 °C and can achieve an electrical conversion efficiency of 47%. SOFC cells generally operate at a higher temperature, 750–1000 °C, and have shown conversion efficiencies up to 43%. While both can be exploited for CHP applications, the MCFC is a relatively complex cell and has been developed for electricity production alone. The SOFC is simpler and some small SOFC units are being developed for domestic CHP applications. Capacities can range from as low as a few tens of watts to several megawatts.

Costs for all types of fuel cells remain high compared to other sources of electricity generation, but potential for cost reduction exists, particularly for PEM cells, which are being actively developed as power units for electric vehicles.

NUCLEAR POWER

Nuclear power plants exploit the controlled use of the nuclear fission reactions of large periodic table elements such as uranium and plutonium—reactions that release a massive amount of heat energy—to provide heat to generate steam for electrical power generation. In a nuclear power plant most of the available heat is captured and used to generate steam, which drives a condensing steam turbine so there will be little heat available for CHP applications. However, it is possible to adapt a nuclear plant so that some of the heat is available.

Nuclear power has been used in Russia and some other eastern European countries for district heating and for seawater desalination, a form of CHP when combined with power generation. However, nuclear CHP technology has never been adopted in the developed world.

The size and cost of building a nuclear power plant means that nuclear CHP plants, where they are built, are constructed for strategic reasons and form part of a national planning strategy. Small nuclear plants capable of CHP use have been proposed in the past but none has reached commercial maturity. Recent safety concerns have also made the future of nuclear generation appear fragile. Nuclear CHP is, therefore, likely to remain the preserve of countries like Russia and is unlikely to find widespread application for the foreseeable future.

COST OF CHP

The economics of CHP depend on the cost of installing the plant with its heat-recovery and generation system, the operating expenses, and the value of the electricity and the heat that it produces. The value of the electricity can be found

TABLE 6.2 Some Representative CHP System Costs[1,2]

System Type	Estimated Lifetime	Cost
Small natural gas–fired reciprocating engine (< 15 kW)	5–6 years	£3290/kW ($5000/kW)
Large natural gas–fired reciprocating engine (110 kW electrical output)	10–20 years	£890/kW ($1300/kW)
Gas turbine CHP (100 kW electrical output)	10–15 years	£900/kW ($1400/kW)
Steam turbine CHP	20 years	> $2000/kW
Phosphoric acid fuel cell (400 kW electrical output)	10–20 years	$2500/kW
Micro-turbines	10–20 years	$3000–4000/kW

[1] *Peter Mayer,* CHP, *Building LifePlans, Greenspec.*
[2] *Paul Breeze,* The Future of Distributed Generation *and* The Future of Fuel Cells, *Datamonitor, 2009 and 2012.*

Source: Greenspec, Datamonitor.

by calculating the cost incurred by purchasing a similar quantity of electricity from the local supplier. The value of the heat is more difficult to estimate but can usually be found by estimating the cost of fuel that would have to be purchased to generate the heat from an alternative source.

Economics will also depend on the type of consumer that the unit is supplying. Large industrial consumers can buy electricity at wholesale prices from suppliers. The cost of the same electricity when supplied to a domestic consumer can be several times that of the wholesale price. Therefore, a domestic CHP system can be relatively expensive and yet still economically viable.

Table 6.2 contains some representative capital costs for CHP systems at the end of the first decade of the 21st century. Some are in pounds sterling and some in U.S. dollars.[5] A small reciprocating engine-based system of less than 15 kW is the most expensive in the table with a cost of around £3290/kW or close to $5000/kW. Small CHP gas engines tend to be expensive but the cost falls as the size increases so that a similar system, but with an electrical generating capacity of 110 kW instead of less than 15 kW, has an estimated cost of £890/kW ($1300/kW). This is comparable to the cost of a similarly sized gas turbine–based CHP system that, according to the data in Table 6.2, would cost £900/kW ($1400/kW). These two large CHP systems would be expected to have a

5. A reasonable but rough conversion can be carried out putting one pound sterling equivalent to one and one-half U.S. dollars.

lifetime up to 20 years but the small reciprocating engine operating at around 5000 hours each year or more would have a lifetime of only 5–6 years.

Other systems included in Table 6.2 are a steam turbine–based system with a cost in excess of \$2000/kW. This price is for a large steam turbine–based plant and the cost will rise for smaller systems. A commercial fuel cell based on PAFC technology is available in the United States for around \$2500/kW. Meanwhile, micro-turbines are still relatively expensive at \$3000–4000/kW.

To obtain the best economic return it is important to size a CHP system correctly. The unit needs to operate for at least 4000 hours each year to be cost effective, particularly where smaller systems are under consideration. Oversized engines running at less than full output will often lose some economic benefit, particularly if this involves dumping heat because it is not required.

The figures in Table 6.2 are all for generation systems based on combustion of natural gas. One of the major renewable CHP options involves wood-burning CHP. There are various approaches to this including a straightforward wood-fired boiler generating steam for a steam generator and wood gasification. The typical price for a relatively small wood-based combustion system (500 kW) is around £2000/kW (\$3000/kW). Larger systems can cost less, with prices as low as £1200/kW (\$1800/kW) feasible.[6]

6. These figures are from Woodfuel Wales.

Chapter 7

Fuel Cells

The fuel cell is an electrochemical device, closely related to the battery, which harnesses a chemical reaction between two reactants to produce electricity. A battery is usually intended as a portable or self-contained source of electricity, and it must carry the reactants it needs to generate electricity within it. Once they are exhausted the battery can no longer supply any power. A fuel cell, by contrast, does not contain any chemical reactants itself but is supplied with them from an external source. So long as these reactants are made available, the cell will continue to provide power.

Batteries come in many different forms, each exploiting reactions between a variety of different chemicals to generate electricity. Fuel cells are unique in that almost all of them exploit a single reaction—that between hydrogen and oxygen to produce water. (There is one exception, the direct methanol fuel cell, which utilizes methanol and oxygen and produces carbon dioxide and water as its by-products.) The simplicity of the basic fuel cell reaction makes it both technically and environmentally attractive—the latter since the only by-products are water and some heat. However, in practice, fuel cells are not quite so simple.

The oxygen for a fuel cell can be supplied from air and air is normally provided to one of the cells electrodes. However, there is no ready source of hydrogen available today. To overcome this, most fuel cells utilize natural gas that they convert to hydrogen in a process known as reforming. This generates carbon dioxide as a by-product, somewhat tarnishing the otherwise spotless environmental credentials of the fuel cell. Even so fuel cells are relatively benign, environmentally, compared to other fossil fuel–generating technologies.

Aside from its environmental performance, one of the most attractive features of the fuel cell is the fact that its energy conversion process is not limited by Carnot cycle thermodynamics that restrict the ultimate efficiency at which any type of heat engine can operate. The best practical simple cycle heat engine efficiency is 50%, achieved in a diesel engine. By contrast, the best fuel cell—an alkaline fuel cell provided with hydrogen reactant—can achieve 60% efficiency today.

The theoretical maximum efficiency for a hydrogen fuel cell operating at room temperature is 83%, but in practice most fuel cells operate at elevated temperatures, reducing efficiency. However, some high-temperature fuel cells can operate in a variety of combined cycle configurations allowing additional gains in efficiency. Meanwhile, overall cycle efficiency is reduced when natural gas must be reformed since this is an energy-intensive reaction. Even so, practical cells are competitive on an efficiency basis with other fossil fuel technologies. If a hydrogen economy ever develops, with hydrogen replacing natural gas as the primary gaseous fuel, then fuel cells should be able to outperform the competing combustion technologies with ease.

The fuel cell has other advantages too. The cell itself has no moving parts and can operate for long periods without maintenance—far longer than any turbine or engine-based generating system. The absence of moving parts makes them inherently quiet too (although this is limited by the use of mechanical pumps, which generate noise), and they emit relatively low levels of pollution compared to other types of generating system based on fossil fuel.

With so much going for them, why are there no fleets of fuel cell power plants today? The answer is cost. While the fuel cell principle has been known since the first half of the 19th century, development of a cheap version of the device has proved extremely challenging. As a result, the first commercial fuel cells only appeared in the early 1990s. Since then research into a wide range of fuel cells has advanced and a variety of commercial units are available. Cost still remains a challenge, but interest in fuel cells as potential power units for vehicles has spurred research that could see some types eventually achieve competitively low prices. Meanwhile, domestic heat and power units based on fuel cells are opening potential new markets for these devices.

HISTORY OF FUEL CELLS

If a voltage is applied to water by placing two electrodes into the liquid and connecting a battery across them, the voltage induces a chemical reaction; hydrogen is produced at one electrode and oxygen at the other. Electrical energy is being used to force the reaction that causes water to dissociate into hydrogen and oxygen.

In 1839, Sir William Grove observed that this process, known as electrolysis, will proceed in the other direction too, and of its own accord. If two specially selected electrodes are placed in water and gaseous hydrogen and oxygen provided (one gas at each electrode), hydrogen will react at one electrode and oxygen at the other, producing water and an electrical voltage between the electrodes. This is the basis of fuel cell operation.

Although the principle was known in 1839, it was not until a century later that the English scientist Francis Bacon was able to develop a practical fuel cell. He demonstrated a 5 kW fuel cell stack in 1959 that he patented as the Bacon Cell. In the same year, Harry Ihrig, an engineer working for the Allis Chalmers Manufacturing Company, demonstrated a 20 horsepower tractor powered by fuel cells. Both were based on the alkaline fuel cell.

By the early 1960s, the Pratt and Whitney Aircraft Corporation had licensed Bacon's fuel cell and began to develop the technology. At around the same time, scientists at U.S. company General Electric began to develop fuel cells based on proton exchange membranes. When NASA was planning its manned space program it identified fuel cells as the ideal power source for its space vehicles. Initially the organization chose the General Electric design and this was used in the Gemini space program. Subsequently it switched to the alkaline fuel cell for the Apollo program and this continued to be used in the space shuttle until it was decommissioned.

Interest in these and other fuel cells continued in the 1960s, 1970s, and 1980s, and a variety of new types of fuel cell were pursued. This led, in 1992, to the launch of the first commercial fuel cell by a company called International Fuel Cells, the parent company of which, United Technologies, owned Pratt and Whitney. This new fuel cell was based on a new acidic electrolyte called phosphoric acid, and the phosphoric acid fuel cell has since developed into the most successful commercial type of cell.

Meanwhile, environmental concerns that began to arise in the 1980s together with fears about oil supply led at the beginning of the 1990s to a revival of an old quest to develop practical electric vehicles that could replace the petrol engine car. While several potential technologies have been pursued, one of the leading contenders is a fuel cell power unit based on proton exchange membrane fuel cells. Development of these continues to receive high levels of investment.

Other types of fuel cell have also been developed. Among these are two high-temperature fuel cells: the molten carbonate fuel cell and the solid oxide fuel cell. The latter is an all solid-state fuel cell, making it potentially extremely durable. Alkaline fuels remain under development too and a sixth type, the direct methanol fuel cell, is seen as a potential new portable source of power.

FUEL CELL PRINCIPLE

The fuel cell principle belongs to a branch of chemistry called electrochemistry. This explores how electricity can be derived from a chemical reaction. In nature certain materials will react with one another spontaneously if the conditions are correct. For example, a strong acid like sulfuric acid will react vigorously with a variety of different materials if they are mixed with it. Other reactions will also proceed spontaneously but require a kick-start. Typical of these is the reaction of hydrogen with oxygen. The two gases can be mixed at ambient temperatures without any reaction, but once the temperature is raised beyond a certain point, with a spark for example, the reaction will proceed explosively.

Spontaneous reactions of this type are called exothermic reactions because they all release heat energy when they are allowed to proceed freely. All the chemical reactions that can be exploited to generate electrical energy are spontaneous reactions. However, when the reaction is exploited electrochemically, the progress of the reaction is managed in such a way that some of the energy that would normally be released as heat emerges instead as electrical energy.

The reaction between hydrogen and oxygen is exothermic. This reaction can be expressed by a simple chemical formula:

$$2H_2 + O_2 = 2H_2O$$

The formula shows two hydrogen molecules and one oxygen molecule reacting to create two molecules of water. Although the formula looks simple, this is in fact a relatively complex process, but it can be broken down into three simple partial reactions, each of which must take place for the reaction to run to completion. The first of these involves the hydrogen molecule, H_2, splitting into two hydrogen atoms, H, and each of these releasing an electron to form a positively charged hydrogen ion, a proton:

$$H_2 = 2H^+ + 2e^-$$

A parallel second partial reaction involves the oxygen molecule, O_2, which splits into two oxygen atoms, O. Each of these absorbs two electrons released from two hydrogen atoms to produce a doubly negatively charged oxygen ion, O^{2-}:

$$O_2 + 4e^- = 2O^{2-}$$

The third and final part of the reaction involves the negatively charged oxygen ion attracting two positively charged hydrogen atoms and the three ions coalescing to form a water molecule, H_2O:

$$O^{2-} + 2H^+ = H_2O$$

Then the reaction is complete.

When hydrogen burns in air, the various steps of the reaction occur in the same place at the same time. However, in a fuel cell the hydrogen and oxygen are not allowed to mix. Instead, the reacting gases are introduced separately with hydrogen supplied to one electrode of the cell and oxygen to the other. The two electrodes are separated by a material called the electrolyte.

The electrolyte is the key element in any electrochemical cell (Figure 7.1) because it acts like a filter to both stop the cell reactants mixing directly with one another and to control how the charged ions created during the partial cell reactions are allowed to reach each other. The electrolyte in a fuel cell is impermeable to the gases, hydrogen and oxygen. It will not conduct electricity in the form of electrons either, nor (provided it is an acidic electrolyte) will it conduct the negatively charged oxygen ions. What it will do is conduct positively charged hydrogen ions.

At the hydrogen electrode (called the anode) hydrogen molecules first adhere to the electrode material and then separate into atoms, each subsequently releasing an electron to form a positively charged ion as shown in the preceding equations. In this ionic form the hydrogen can cross the electrolyte boundary and reach the oxygen at the second electrode. The electrons, however, are left behind at the electrode.

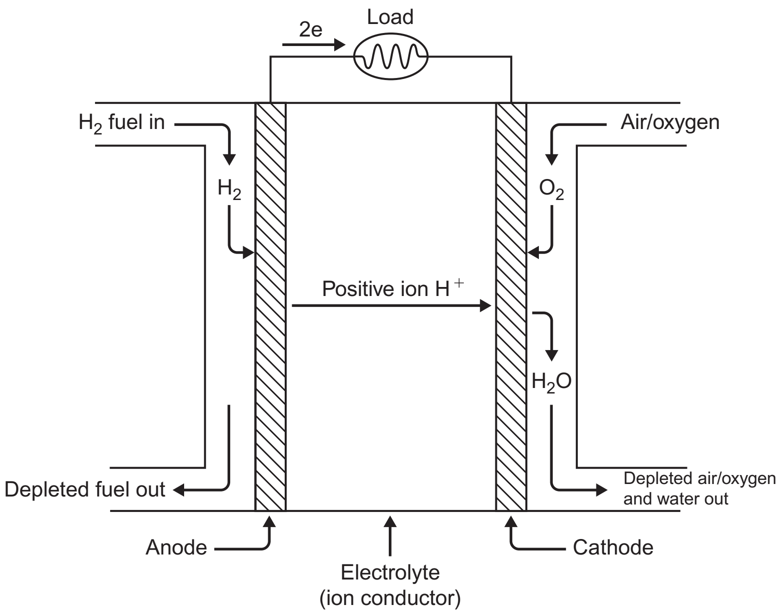

FIGURE 7.1 The principle of the fuel cell.

At the second electrode (called the cathode) oxygen atoms will also adhere to the electrode surface and dissociate, each leaving two oxygen atoms. However, these require a supply of electrons if they are to form oxygen ions. Only in this form can they coalesce with the hydrogen ions traveling through the electrolyte from the anode to create water molecules. The electrons must come from the anode, but these electrons cannot pass through the electrolyte.

If an electrical connection is made between the two electrodes the electrons can now pass through the connecting wire to the oxygen atoms where they will create oxygen ions, allowing the reaction to run to completion. When a small electric light bulb is inserted into this circuit it will glow, proving that an electric current is indeed flowing between one electrode and the other.

All electrochemical cells operate in this way. While part of the reaction can take place within the confines of the cell, it can only be completed if electrons are allowed to travel from one electrode to the other through an external circuit.[1] If no connection is made the exothermic nature of the reaction means that a charge builds up at each electrode, creating an electrical potential, the cell voltage, which will drive electrons from one electrode to the other when a connection is made.

Since the electrolyte is such an important part of an electrochemical cell, fuel cells are generally identified by the type of electrolyte they employ.

1. In fact, a few electrons can pass through the electrolyte, creating a short-circuit current, but its resistance to this is very high.

Phosphoric acid fuel cells use phosphoric acid, which being acidic is a proton (positively charged hydrogen atom) conductor. Proton exchange membrane cells also rely on an acidic membrane to allow hydrogen ions to pass. The alkaline fuel cell is a hydroxide ion (OH^-) conductor; the hydroxide ion is the complement in water of the hydrogen ion. Meanwhile, the high-temperature solid oxide fuel cell uses a solid that is a conductor of oxygen ions. Other fuel cell reactions appear more complex even though the end result is the same. So, for example, the molten carbonate fuel cell electrolyte conducts carbonate ions.

CATALYSTS

The preceding description of the operation of a fuel cell is a simplification because it omits one key feature of the reaction between hydrogen and oxygen. Although hydrogen atoms and oxygen atoms will react spontaneously to form water, both hydrogen and oxygen are found (at room temperature) in the molecular forms H_2 and O_2. These hydrogen and oxygen molecules must split into atoms before the reaction will proceed, but they will not do so spontaneously because of the chemical bond holding each molecule together, even though these bonds are much weaker than the chemical bonds that will bind them together into a water molecule. The energy needed to cause the individual molecules to dissociate into atoms creates an energy barrier called the activation energy that must be overcome before the highly exothermic reaction between the atoms can take place.

One method of splitting the molecules is to raise their temperature. They will start to dissociate rapidly between 800 °C and 1000 °C. A flame or a spark will be hot enough to split sufficient of the molecules to start the reaction, which then generates so much heat spontaneously that it keeps the reaction going until all the hydrogen or oxygen is used up. Some fuel cell designs use high temperatures to encourage the gas molecules to dissociate, but high temperatures bring their own design and materials problems.

The alternative is to use a catalyst. A catalyst is a component that is needed for a reaction to take place but is not actually consumed during the reaction. As such, it is usually used to accelerate a reaction that will otherwise take place slowly. In the case of the fuel cell the best catalyst for low-temperature acceleration of the reaction is metallic platinum. The platinum acts by attracting the hydrogen and oxygen molecules that will preferentially stick to its surface in dissociated form, generating a supply of the atoms needed for the fuel cell reaction to take place. In its presence, the reaction can take place below 100 °C.

Platinum, even though it is only required in small quantities in a fuel cell, is expensive and this helps to elevate the cost of the cells themselves. A key area of fuel cell research is therefore directed at finding cheaper alternatives. Platinum is also very sensitive to impurities in the gaseous reactants that can poison it, rendering it ineffective. Sulfur dioxide is a particular problem and so is carbon monoxide, both of which can find their way into hydrogen generated by the reforming of natural gas or other fuels. This is another reason why alternative catalysts, which are less sensitive to poisoning, are being sought.

HYDROCARBON GAS REFORMATION

A fuel cell is designed to "burn" hydrogen and oxygen to generate electricity. Hydrogen is not generally available, but hydrocarbon gases such as natural gas or even gases generated from biomass can be converted into hydrogen in a process known as reformation. The reformation reaction generates a gas that contains a mixture of carbon dioxide and hydrogen that can then be supplied as a reactant for the fuel cell. (The carbon dioxide will be inert and so will not interfere with the reaction other than by acting as a diluent for the hydrogen.) The main constituent of natural gas and of most biogas is methane, and this is the main target for the standard reformation process, although other hydrocarbons can also be reformed and even coal can be converted into a hydrogen-rich fuel if necessary (see coal gasification in Chapter 3).

The conversion is usually carried out as a two-stage process. In the first stage the methane-rich gas is mixed with water vapor and passed over a catalyst at a high temperature where the gases react to produce a mixture of hydrogen and carbon monoxide, a process called steam reforming. A second reaction, called the water shift reaction, is then carried out during which additional water vapor is added to the new mixture where it reacts with the carbon monoxide, again in the presence of a catalyst, to produce more hydrogen and carbon dioxide.

The degree to which this second stage reaction is carried to completion is extremely important for fuel cells because the catalysts in low-temperature cells are sensitive to carbon monoxide poisoning. In consequence, virtually all the carbon monoxide must be removed from the fuel before it is fed into the fuel cell. Natural gas can also contain some sulfur impurity. This is normally removed by cleaning before the fuel is reformed, but if not, then any remaining sulfur must be removed since low-temperature fuel cell catalysts are extremely sensitive to its presence.

While natural gas is the most convenient source of hydrogen for a fuel cell today, other fuels can also be exploited. For example, methanol can also be converted in to a hydrogen-rich gas using a reforming process, as can gasoline, though the latter requires an extremely high temperature (800 °C). Both these processes are of interest to the automotive industry.

Since reforming of all these fuels takes place at a relatively high temperature, low-temperature fuel cells have to be equipped with an external reformer to process the fuel before it enters the fuel cell. However, heat generated within the fuel cell may be used to help drive the reforming process. The conditions inside the two main types of high-temperature fuel cells are sufficient for the reforming to take place within the cell itself, simplifying system design.

While a fuel cell burning hydrogen and oxygen produces no carbon dioxide, most fuel cells will generate carbon dioxide because they derive their hydrogen from natural gas or another carbon-containing fuel. When methane is converted into hydrogen it generates exactly the same amount of carbon monoxide as it would have generated if it had been burned in a gas turbine. However, if hydrogen can be generated without the need for fossil fuel combustion—by using hydropower to electrolyze water, for example—then burning the gas in a fuel cell provides an efficient and emission-free source of electricity.

FUEL CELL EFFICIENCY

The reaction between hydrogen and oxygen to create water releases a precisely quantified amount of energy. Not all this energy can be converted into electricity because some is required to overcome the energy barrier that normally prevents the reaction proceeding. When this reaction takes place between oxygen and hydrogen, each provided at a pressure of one atmosphere and at room temperature, the theoretical maximum chemical-to-electrical energy conversion efficiency that can be achieved is 83%.

This efficiency is an ideal and no cell would be able to achieve that figure. The precise reaction conditions will affect the potential efficiency too. If the pressure of the reacting gases is increased, then conversion efficiency can be increased. On the other hand, increasing the operational temperature of the cell will reduce the overall efficiency that can be achieved.

The actual conversion efficiency of a fuel cell is reflected in the voltage that the cell produces between its terminals. The theoretical maximum cell voltage at open circuit for a fuel cell operating at room temperature, when the cell is delivering no current, is 1.229 V. In practice, such a cell would deliver a voltage less than this because of losses resulting from the internal resistance of the cell and activation energy barriers of various sorts at electrode interfaces.

This theoretical maximum voltage only applies to a cell where the cell reaction product is water in liquid form. In most practical cells, where the product is actually water vapor, this maximum falls to 1.18 V. At 100 °C this falls to 1.16 V, and at 800 °C the ideal cell voltage is only 0.99 V. This is equivalent to a maximum ideal efficiency of 67%. In practice, high-temperature cells might approach 50% efficiency. Low-temperature cells can do better than this.

While high-temperature cells are ostensibly less efficient, the loss of electrochemical efficiency is not necessarily a major handicap. The heat generated within the cells can be exploited either to produce more electricity in some form of hybrid system, or it may be utilized in a combined heat and power system, providing useful heat as well as electricity.

FUEL CELL TYPES

There are six principal types of fuel cell that are of interest to the power generation market. They are listed in Table 7.1 together with their primary characteristics.

TABLE 7.1 Fuel Cell Characteristics

	Operating Temperature (°C)	Catalyst	Efficiency (%)
Alkaline fuel cell	25–200	Noble metal catalyst	60
Phosphoric acid fuel cell	150–200	Platinum	35–42
Proton exchange membrane fuel cell	80	Platinum	30–50
Molten carbonate fuel cell	650	Nickel	47
Solid oxide fuel cell	750–1000	None needed at temperatures used	35–43
Direct methanol fuel cell	60–130	Platinum/ruthenium	25–40

The alkaline fuel cell is the earliest fuel cell to be developed and it still remains under active development. Early AFCs operated between 150 °C and 200 °C, but more recent designs allow the cell to operate at a much lower temperature of 25–70 °C. The cell electrolyte is a solution of potassium hydroxide. AFCs have achieved the highest efficiency of any fuel cell with practical efficiencies of 60% in the space shuttle and experimental efficiencies of 70%.

The phosphoric acid fuel cell was the first to achieve commercial success for mainstream power generation. It too is a low-temperature cell with an operating temperature of 150–200 °C and an acidic electrolyte composed of phosphoric acid. Practical cell efficiencies of 42% have been recorded. The proton exchange membrane fuel cell is being developed as a potential automotive power source. Its electrolyte is primarily water so standard cells must operate below 100 °C. Efficiencies up to 50% are achievable in practice and up to 60% may be possible with a pure hydrogen fuel source.

Two high-temperature fuel cells have also reached a high stage of maturity. The molten carbonate fuel cell operates around 650 °C and can reach 47% efficiency, while the solid oxide fuel cell, with an operating temperature range of 750–1000 °C depending on the electrolyte, can achieve 43% efficiency in large units. Both high-temperature cells can also be incorporated into hybrid cycles to increase overall generation efficiency.

The sixth type of cell is the direct methanol fuel cell. This is another type of polymer membrane cell with an operating temperature up to 130 °C. The fuel fed to the anode is methanol mixed with water, which can be consumed directly without reforming. The fact that the cell can operate with a liquid fuel makes it attractive both for automotive applications and for portable equipment such as computers and mobile phones. Practical efficiencies are low at around 25%, although experimental cells have reached 40%.

Alkaline Fuel Cell

The alkaline fuel cell (AFC) employs an electrolyte composed of a concentrated solution of potassium hydroxide (Figure 7.2). This is the alkali that has the highest conductivity of the alkaline hydroxides. The earliest practical cell of this type, developed by Francis Bacon, used a concentration of 45%. The cell requires a metal catalyst such as platinum, but the actual metal used varies with application.

The cell reaction in the AFC is slightly different to the standard fuel cell reaction because the electrolyte conducts hydroxide ions (OH^-) rather than protons. Molecular hydrogen is supplied to the anode where it splits into hydrogen atoms that each release an electron to form a hydrogen ion. Meanwhile, at the cathode oxygen molecules dissociate to form atoms, but in this case they take up two electrons and then react with water to form hydroxide ions. These can then pass through the electrolyte to the anode where they react with the hydrogen ions generated there to regenerate water. The series of reactions are as follows:

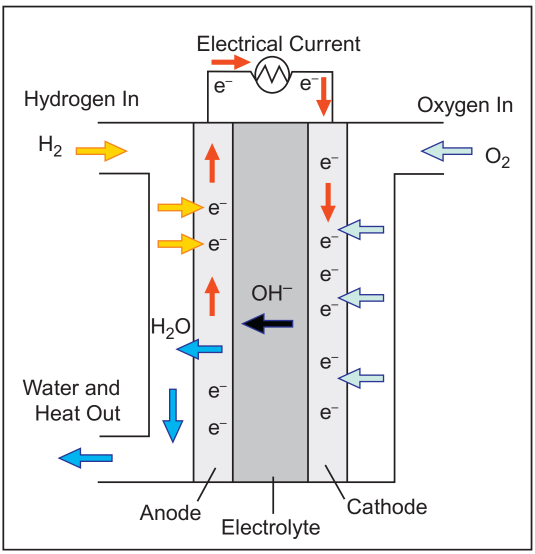

FIGURE 7.2 Alkaline fuel cell.

$$H_2 = 2H^+ + 2e^-$$
$$O_2 + 4e^- = 2O^{2-}$$
$$O^{2-} + H_2O = 2OH^-$$
$$2H^+ + 2OH^- = 2H_2O$$

In essence, the reaction is exactly the same as the standard fuel cell reaction described before but it is mediated by the hydroxide ions. These are generated at the cathode in the reaction between oxygen ions and water and then consumed at the cathode.

Early AFC cells operated at a relatively high temperature around 260 °C and some modern cells use similar temperatures. Other modern cells utilize much lower temperatures with the cells in the space shuttle operating at 85–95 °C. These cells are capable of practical efficiencies of 60%. Developers are hoping to develop cells that can operate at ambient temperature.

In a classic AFC cell used in the space program the electrolyte is held within a porous solid matrix of an asbestos-like material. Electrodes are generally exotic to achieve the high efficiencies required for such applications. Space shuttle cell anodes were 80% platinum and 20% palladium carried on a silver-plated nickel mesh and cathodes 90% gold and 10% palladium on a similar mesh. For less demanding applications simpler catalysts can be employed.

Cell lifetime has so far been the main factor limiting the use of the AFC. The lifetime of a space shuttle fuel cell was around 2600 hours. Modern developments have begun to extend this and a lifetime of 5000 hours appears feasible over the short term. Limiting factors include the buildup of potassium carbonate generated from carbon dioxide both in air and as an impurity within the hydrogen fuel.

As will other fuel cells, the AFC requires hydrogen and oxygen. The space program cells used high-purity hydrogen and oxygen as fuel, but for terrestrial applications air is preferred for the cathode. This must be treated first to remove all traces of carbon dioxide. There are sources of pure hydrogen such as hydrazine, but for practical applications reformed natural gas is the best source. However, this also must be purged of carbon dioxide to maintain cell health.

One of the key recent AFC developments is the use of a circulating electrolyte rather than an electrolyte immobilized within a matrix. Cycling the electrolyte allows carbonate impurities to be removed, potentially increasing cell lifetime. Other research is aimed at developing hydroxyl ion exchange membranes similar to the proton exchange membranes already developed for acidic fuel cells.

There are a small number of companies currently developing AFCs and the availability of commercial cells is limited. They are only used today where efficiency and reliability are more important than cost. Outside the space program this has been restricted to a number of vehicle applications such as fork-lift trucks and golf trolleys. Costs are too high and lifetimes too short for their use for general automotive applications, but if both could be reduced they would have great potential.

Phosphoric Acid Fuel Cell

The phosphoric acid fuel cell (PAFC) uses an electrolyte composed of pure phosphoric acid, H_3PO_4 (Figure 7.3). This acid, which is a solid up to 42 °C, is a relatively poor conductor of hydrogen ions but is stable in liquid form to just above 200 °C, so it can form the basis of a fuel cell that operates using hydrogen and oxygen supplied as air. Most PAFCs operate between 150 °C and 200 °C.

The electrolyte, once liquefied, is contained in a silicon carbide matrix where it is held within the pores of the material by capillary action. Since the electrolyte is a liquid, care must be taken to control evaporation or migration as this will impair the operation of the cell. PAFCs are generally operated at atmospheric pressure but higher-pressure operation is possible. This will increase efficiency but may also lead to higher cell corrosion rates by increasing the reactivity of the acidic electrolyte.

Electrodes for the cell are made from porous carbon paper that is heat-treated to create a large surface area and then coated with a fine layer of platinum or platinum alloy onto which the gaseous reactants can be absorbed. These electrodes are bonded to the electrolyte-supporting matrix using Teflon mounts. There are grooves in the back of the electrodes that carry the hydrogen or oxygen to each cell. Carbon is a good electrical conductor, so it can also be used to transport the current from the cells. Each cell produces a voltage around 0.65 V. Cells can be connected back-to-back, in series, to build what is known as a stack. Stacks are then connected in parallel to provide the required current and voltage output.

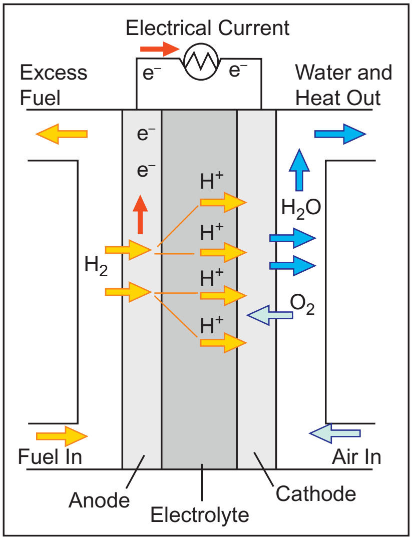

FIGURE 7.3 Phosphoric acid fuel cell.

The PAFC requires hydrogen and oxygen, the latter from air. The cell must be heated to its operating temperature before it can be started, but once it starts operating the cell reaction produces sufficient heat to maintain its temperature. Water generated at the cathode by the cell reaction emerges as water vapor, which is swept away from the cell in the airstream feeding oxygen to the cathode. It is important that water removal is effective. Otherwise, the water might dissolve in the electrolyte.

Most modern cells operate at the highest temperature that is practical, close to 200 °C, to reduce the sensitivity of the catalyst to poisoning and to ensure that all the water from the cell reaction is produced in vapor form. Above this temperature the electrolyte starts to decompose and becomes more aggressive in its reactivity, particularly with the carbon components of the cell. At the cell operating temperature the hydrogen–oxygen reaction will proceed sufficiently swiftly with a platinum catalyst to sustain cell operation. However, the concentration of carbon monoxide in the hydrogen must be kept below 1.5% to prevent catalyst poisoning. Sulfur tolerance is virtually zero, so this must be scrupulously removed too.

Hydrogen for cell operation is normally provided by reforming natural gas and packaged PAFC units include a reformer. Heat generated within the cell during operation can be used to help drive this reaction. Other fuel sources are possible too, including biogas. The reforming of natural gas or biogas produces carbon dioxide, but the low-temperature reforming process does not produce significant amounts of nitrogen oxide and other emission levels are low.

Practical cell efficiencies between 36% and 42% are typical, with higher efficiencies generally achieved with pressurized operation. The operational temperature makes PAFC units suitable for combined heat and power (CHP) generation, producing hot water. Packaged CHP units can reach up to 87% energy utilization. PAFC lifetimes up to 40,000 hours have been recorded, providing good operational potential, and developers plan to extend this further. They have been operated at temperatures as low as −32 °C and as high as 87 °C.

Most of the development of PAFCs has taken place either in the United States or Japan. The U.S. company United Technologies Corporation (UTC) launched the first commercial units in 1992 with an electrical output of 200 kW. The newest version of this cell is now available in 200 kW and 400 kW unit sizes. Both can also provide hot water. UTC has worked in partnership with the Japanese company Toshiba. Fuji Electric Corporation and Mitsubishi Electric have also been active in PAFC development, and the former is marketing a 100 kW unit. In all cases larger capacity can be achieved with multiple units.

The applications of PAFC units are varied but many have been installed to provide power and heat in institutions, such as hospitals, at commercial sites, and in office blocks. They are also suitable for backup use. Early cells were expensive and they appeared to have little future at the beginning of the 21st century as other types of cells reached maturity. However, the development of those other cells has been much slower than expected. In the meantime the continued development of PAFC units and an increase in their market size, leading to economies of scale, has meant that prices have fallen. At the end of the first decade of the 21st century they were the most successful of all fuel cells for power generation applications.

Proton Exchange Membrane Fuel Cell

The proton exchange membrane (PEM) fuel cell uses a polymer membrane as its electrolyte. The cell was invented by U.S. company General Electric and tested for U.S. military use in the early 1960s and was adopted by NASA for its Gemini space program. After development for the U.S. Navy it was adopted by the British Navy in the early 1980s. Transportation applications have continued with the Canadian company Ballard Systems developing units for buses. The low weight of the cell and its high efficiency have made it a major candidate for future automotive applications, which has attracted large investment from many automotive companies. Power generation applications have also developed alongside these other applications.

The membrane that forms the electrolyte of the PEM cell is usually a compound called poly-perfluorocarbon sulfonate. This is a close relative of Teflon but with acidic sulfonate molecular groups attached to its polymer backbone to provide conductivity. In its normal state the membrane is not conductive, but if it is allowed to become saturated with water the acidic groups attached to the membrane release protons, allowing it to conduct hydrogen ions. The membrane itself is usually between 50 microns and 175 microns thick, the latter equivalent to seven sheets of paper. As a consequence of the fact that the conductivity is provided by water, the cell must be kept well below the boiling point if it is to remain stable. Practical cells normally operate around 80 °C. This makes its catalyst more susceptible to poisoning than other, higher-temperature cells.

The electrodes of a PEM fuel cell are made from porous carbon-containing platinum that can be printed directly onto the membrane. A further porous carbon-backing layer provides structural strength to each cell as well as supplying electrical connections (Figure 7.4). Typical cell voltage is 0.7–0.8 V. As with the PAFC, cells are joined in series and in parallel to provide sufficient current and voltage.

A PEM fuel cell requires both hydrogen and oxygen to operate. Oxygen is normally supplied from air and the hydrogen by reforming natural gas. This hydrogen must be purged of carbon monoxide and sulfur compounds to avoid catalyst poisoning. The low cell operating temperature also means that there is no heat from the cell to use to help drive the reforming reaction so additional natural gas must be burned, reducing overall cell efficiency.

When supplied with pure hydrogen, a PEM fuel cell is capable of a theoretical fuel-to-electrical conversion efficiency of 60% but in practice the

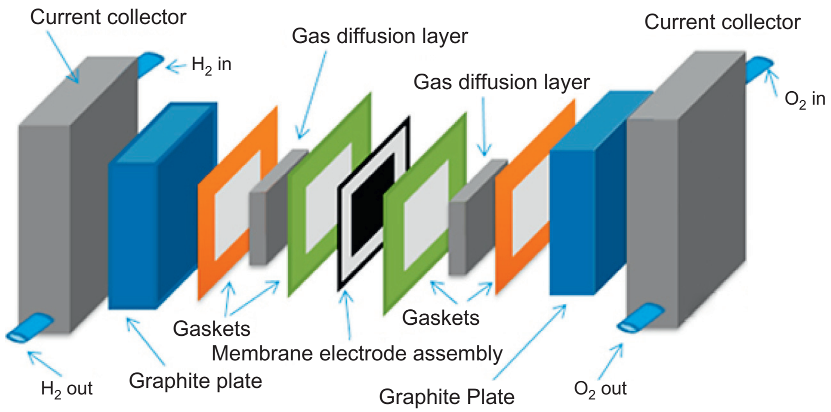

FIGURE 7.4 Proton exchange membrane fuel cell.

efficiency is likely to be closer to 50% (higher efficiency has been recorded in the laboratory). This relatively high efficiency is one of the attractions for the automotive industry. When the hydrogen is derived from natural gas by reforming, efficiency falls so that practical natural gas–based PEM fuel cells can achieve an efficiency around 42%. The need for a reformer also leads to a startup time around 20 minutes while the reformer reaches its operating temperature. With hydrogen, the startup time would be virtually instantaneous.

Efficiency might be improved with higher-temperature operation and this would also reduce sensitivity to catalyst poisoning. This is driving research into newer types of membranes. One that has achieved some success is a polybenzimidazole membrane containing phosphoric acid that can operate between 160 °C and 180 °C, and a number of companies are offering cells that operate in this temperature range.

The operational lifetime of a PEM fuel cell has been lower than that of a PAFC with practical lifetimes of 10,000 hours typical. This is considered adequate for automotive applications but is too short for many stationary applications. However, some products claim a lifetime of 40,000 hours.

There are a range of stationary applications for PEM fuel cells but they have yet to establish themselves in any one part of the market. Early development was directed at large stationary cell stacks with capacities of up to 250 kW and there has been one 1 MW utility application. However, the largest units today are likely to be below 100 kW. There has been significant growth at the small end of the market with units up to 10 kW that can be used for backup and standby. Smaller units, perhaps around 3 kW and less, are being developed for the domestic market where they can supply both power and low-grade hot water suitable for household use and heating.

Another application that appears to be growing is for portable power generation. These units take advantage of the high efficiency of the PEM fuel cell when operating with hydrogen to provide a cheap, highly efficient source of portable power. However, this is a market to which the direct methanol fuel cell is also directed.

Much of the pioneering work on the PEM fuel cell was carried out in the United States and Canada but in recent years Japanese companies have also taken an active interest. Japan is also home to an active domestic fuel cell program at which a number of products are directed. There is also interest in Europe, mostly for small PEM fuels for portable and standby power and for domestic use. The main global research drive is to find the means to reduce the amount of platinum needed in a PEM fuel cell. While alternative catalysts may be possible, platinum remains the best material, but it is expense and the cost can vary wildly due to global commodity market volatility.

Molten Carbonate Fuel Cell

The molten carbonate fuel cell (MCFC) has an electrolyte that is composed of a mixture of carbonate salts (Figure 7.5).[2] These are solid at room temperature, but at the cell operating temperature of 650 °C they become liquid. Work on

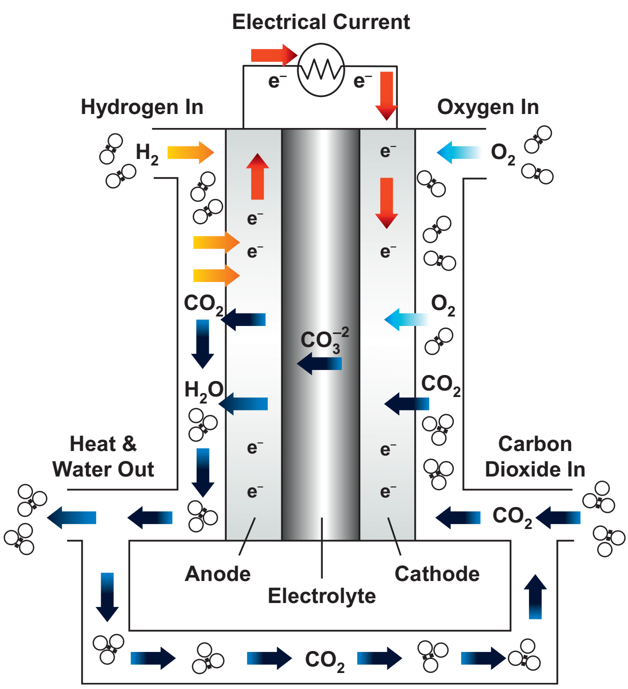

FIGURE 7.5 Molten carbonate fuel cell.

2. The electrolyte is usually a mixture of sodium carbonate, potassium carbonate, and lithium carbonate.

high-temperature MCFCs began during the 1950s after earlier research in the 1930s into solid (low-temperature) carbonate fuel cells made little progress. The U.S. Army tested cells during the 1960s and in the 1970s the U.S. Department of Energy began to support research. Japanese companies also picked up the technology. MCFC fuel cells are more complex than other fuel cells and are most cost effective in large stack sizes of 250 kW or more. At the beginning of the second decade of the 21st century the market for MCFCs was dominated by U.S. company FuelCell Energy (FCE), which has formed a strong partnership with Korean company POSCO Energy.

The MCFC has the most complex fuel cell reaction of all the cells available commercially. The electrolyte is a mixture of alkali metal carbonates that, when heated above 650 °C, become molten and capable of conducting carbonate ions (CO_3^{2-}). The molten carbonate mixture is held within a solid matrix that is commonly made from lithium aluminum oxide. Hydrogen is fed to the anode of the cell where it reacts with carbonate ions in the electrolyte according to the equation:

$$H_2 + CO_3^{2-} = H_2O + CO_2 + 2e^-$$

During this reaction part of the electrolyte is consumed and carbon dioxide is released as a gas, mixing with the hydrogen stream. This carbon dioxide is fed back to the cathode where it reacts with oxygen to regenerate carbonate ions in the electrolyte according to the following equation:

$$O_2 + 2CO_2 + 4e^- = 2CO_3^{2-}$$

The overall reaction when these two electrode reactions are taken together is simply that of hydrogen and oxygen as in other fuel cells, but it is mediated by the carbonate ions in the electrolyte. Carbon monoxide in the hydrogen fuel can react at the anode to generate more hydrogen in a shift reaction. It can also, in principle, react with carbonate ions and so can form a fuel for the cell.

The high operating temperature of the cell means that the cell reaction takes place without the need for an expensive platinum catalyst and cheaper, nickel-based electrodes are normally employed. The anode is generally a metallic nickel alloy while the cathode is nickel oxide. The latter will dissolve slowly in the hot carbonate electrolyte, which can limit cell life. These electrodes are applied to the outer surfaces of the refractory porous tile in which the electrolyte is held. Cell voltage is 0.8 V.

The reason why such a complex cell has proved worth developing lies in the potential efficiency. The theoretical conversion efficiency is 60% although production units may only achieve around 50%. This does not represent the efficiency limit, however. A MCFC fuel cell produces high-temperature waste heat and this can be exploited in a small gas turbine or a micro-turbine, without the need for additional fuel, to generate more electricity. For greatest efficiency this requires the cell to be pressurized, but in this configuration the MCFC fuel cell may be theoretically capable of between 75% and 80% overall efficiency. Alternatively, the high-grade heat produced by the cell can be used for cogeneration.

The relative complexity of the MCFC means that it is not economical to manufacture and operate very small units. Some of the units that have been developed are over 100 kW in capacity and most are two or three times that size.

The company that dominates the MCFC market is FCE, which was established around 1976 and worked with the U.S. Department of Energy and the U.S. Electric Power Research Institute to develop molten carbonate technology. The company has formed commercial links with companies in Europe and Asia including a partnership with South Korean company POSCO Power. The company's current range of units includes modules of 300 kW, 1.4 MW, and 2.8 MW, all rated at 47% efficiency.

Applications for the MCFC are generally for large distributed generation. In the United States FCE has developed units that can be run on biogas. It has also developed units that can provide hydrogen as well as electricity by boosting the reforming of natural gas or biogas within the cell.

Other companies that are involved in the development of MCFC units include Italian company Ansaldo and Japanese company Ishikawajima-Harima Heavy Industries. The latter has been working on MCFCs since 1983 and has been involved in government-sponsored development programs. It has tested 300 kW pressurized MCFC stacks operating at a pressure of 4 bars. In South Korea, Doosan Heavy Industries and Construction has also been involved in MCFC development.

Solid Oxide Fuel Cell

The solid oxide fuel cell (SOFC) is the second major high-temperature fuel cell under development. It is unique among fuel cells for having a solid electrolyte, making it potentially the most robust of all fuel cells (Figure 7.6). The electrolyte is usually made from zirconium oxide, ZrO_2, zirconia. When traces of other oxides such as yttrium, calcium, or magnesium oxide are added to the zirconia, it becomes capable of conducting oxygen ions. However, this conductivity only becomes significant at very high temperatures and so the cell must operate at around 1000 °C.

Solid oxide electrolytes were first studied during the 1930s, with little success. However, work continued during the 1950s and 1960s. The most persistent program was carried out by U.S. company Westinghouse (now owned by Siemens) in conjunction with the U.S. Department of Energy. This finally established the SOFC as a viable proposition but that program appears to have halted. However, a range of other companies have since taken up the technology.

As will all the other fuel cells discussed here, the cell reaction in the SOFC involves that between hydrogen and oxygen producing water. The solid oxide that forms the cell electrolyte is an electrical insulator so electrons cannot pass

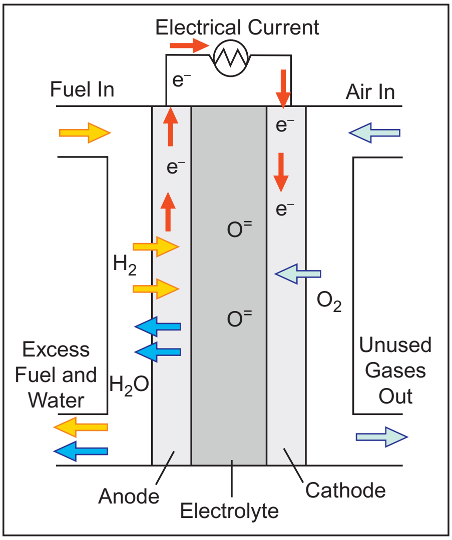

FIGURE 7.6 Solid oxide fuel cell.

through it. Neither will it conduct hydrogen ions. However, at the cell operating temperature the solid structure will allow oxygen ions, O^{2-}, to pass. It is this oxygen ion conductivity that permits the fuel cell to operate. In operation, molecular oxygen is delivered to the cathode of the cell where it dissociates and takes up electrons to produce oxygen ions that migrate through the electrolyte to the anode. Hydrogen is delivered to the anode where it too dissociates, releasing electrons to the external circuit and leaving hydrogen ions that react with the migrating oxygen ions to produce water. At the elevated operating temperature this water is produced as vapor that is swept away in the fuel gas stream.

The electrolyte used in the SOFC is very thin. A thickness of around 100 microns is common but they may be as thin as 10 microns. Electrodes must be bonded to the electrolyte layers and these also serve as a support structure to give the cell strength. Since the cell operating temperature is so high, the different materials used in cell construction must be carefully designed to have the same coefficient of expansion, otherwise the cell would crack apart as it was heated.

The solid electrolyte requires an extremely high operating temperature to facilitate oxygen ion conductivity. The most common zirconia-based electrolytes must be heated above 800 °C to ensure they are adequate for cell operation. Newer materials may be able to provide sufficient conductivity at a much lower temperature. However, the temperature needs to remain above around 600 °C to allow reforming of natural gas to take place within the cell.

The high temperature means that no electrode catalyst is necessary to facilitate generation of hydrogen and oxygen atoms at the electrodes. The anode is normally made from metallic nickel dispersed in a ceramic matrix that is unreactive toward hydrogen and the cathode is made from a conductive oxide that will not react with oxygen. Reformation of natural gas into hydrogen can take place directly on the nickel anode.

There are two primary designs for SOFC. The first was developed by Westinghouse and is based on a tubular cell. Under this design the electrolyte/electrode structure is fabricated as a ceramic tube. The cell cathode is on the inside of the tube and the anode on the outside. This creates an extremely practical way of maintaining separation of hydrogen and oxygen fuel streams since these will easily react at the cell operating temperature. However, the fabrication of the tubular structures is expensive.

The alternative is to use a more conventional planar cell design. This is much cheaper to fabricate but makes the fuel gas routing more complex. Most current SOFC designs use planar cells but the use of tubular cells is being resurrected too.

The theoretical efficiency of an SOFC operating around 1000 °C is 60%. Practical cells have achieved around 43% and the practical limit is likely to be around 50%. However, the high operating temperature of the SOFC, coupled with its robust nature, offers a range of hybrid cycle options that can boost output. A simple addition is a heat-recovery steam generator and steam turbine to create a combined cycle plant analogous to that utilized with gas turbines. This could probably boost efficiency to 60% in large systems. More complex but potentially more efficient would be to operate the SOFC under pressure and then use the hot, high-pressure waste gases to drive a gas turbine (Figure 7.7). Heat could then be recovered in a steam generator and used to drive a steam turbine. In a large plant, this configuration could conceivably achieve 73% efficiency though nothing like that level of efficiency has yet been demonstrated.

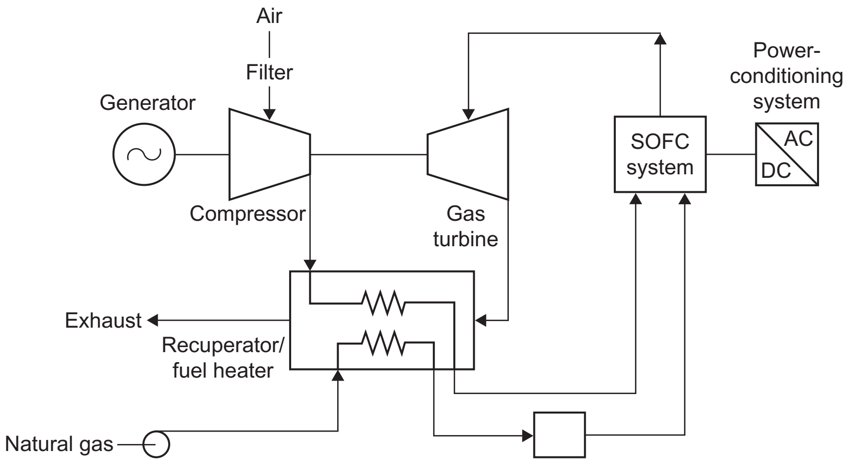

FIGURE 7.7 Block diagram of a SOFC/gas turbine power plant.

The major attraction of the SOFC is the robust nature of its electrolyte that is expected to confer extremely long lifetimes. Units have been tested for 60,000 hours without failure and operating lives of 20 years or more are achievable.

Much of the early development of SOFCs was with the aim of producing large units. Stack sizes were expected to be 100 kW to 250 kW and multimegawatt power plants were proposed. More recently the focus has shifted to much smaller units although U.S. company Bloom Energy is offering 100 kW and 200 kW units that have been popular with data centers, often in multimegawatt installations. Meanwhile, units of 1 kW to 10 kW have been developed for the domestic combined heat and power market, particularly in Germany. Other domestic units are being developed for the Japanese market.

Direct Methanol Fuel Cell

The direct methanol fuel cell (DMFC) is a polymer membrane fuel cell, similar in concept to the PEM fuel cell. The major difference is that in the DMFC the fuel supplied to the anode of the cell is not gaseous hydrogen but methanol in liquid form (Figure 7.8). The methanol, mixed with water, can react directly at the cell electrode without the need for reforming. This simplifies the cell, reducing costs. The use of a liquid rather than a gaseous fuel is extremely attractive too as it makes fuel handling much easier. The main application for the DMFC is as a portable power supply and it is of interest to the automotive industry for the same reason.

Research into the DMFC was carried out in the 1950s and 1960s, then revived during the 1990s. Early cells had exhibited low current densities but

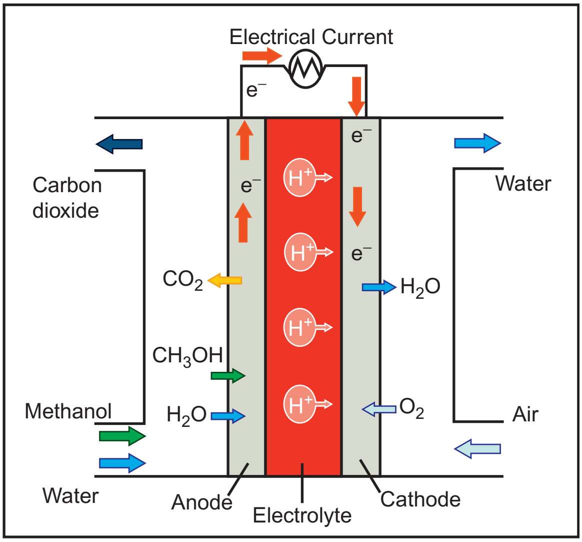

FIGURE 7.8 Direct methanol fuel cell.

this has subsequently been improved. However, efficiency remains poor, with practical cell efficiencies around 25%, low compared to other types of fuel cell. The interest in liquid fuel–based fuel cells has also extended to ethanol and direct ethanol fuel cells are also under development.

The DMFC is novel because its cell reaction combines the reformation of methanol in the presence of water to generate hydrogen at the anode and then the anode reaction of a conventional hydrogen fuel cell. This can be expressed in the two following equations. The first represents the reforming of methanol to generate hydrogen with the production of gaseous carbon dioxide. The second shows the hydrogen atoms that were generated in the first reaction releasing electrons to form hydrogen ions that then migrate through the electrolyte of the cell to the cathode.

$$CH_3OH + H_2O = CO_2 + 6H$$
$$H = H^+ + e^-$$

The cathode of the cell is supplied with oxygen from air and the oxygen subsequently reacts at the electrode to generate water. In a stationary application a DMFC will operate at a relatively high temperature up to 130 °C and the water generated at the cathode will simply be swept away with the air as water vapor. For portable applications, however, it is important to try and recover this water and return it to the anode of the cell for recycling through the cell. This will allow the cell to be supplied with pure methanol fuel. Active water recycling systems are relatively complex, adding to the cost of cells, and research is directed at finding passive methods for recycling water between the cell electrodes.

The membrane electrolyte for the DMFC is an acidic polymer membrane similar to that used in a PEM fuel cell, saturated with water. The anode is made from a mixture of platinum and ruthenium deposited onto a carbon support that can then be printed onto the membrane while the cathode is platinum on carbon. One of the problems with the DMFC is diffusion of methanol through the electrolyte from the anode to the cathode where it can react to produce carbon monoxide, which will poison the cathode catalyst.

A DMFC should be able to achieve an efficiency of 40% but practical cells have not exceeded 25%. This may not be a problem for portable applications provided the cell is cheap enough but would probably be too low for larger stationary applications. Cell lifetimes are problematic too with some early cells capable of only 1000 hours.

In spite of these problems DMFCs are beginning to achieve commercial success as power supplies for portable electronic devices such as laptop computers and mobile phones where they can provide a higher energy density and more power from a single charge of methanol than most batteries. Device stacks are typically 200 W or less. Automotive applications are also attractive but the cells have yet to be proven in this application. It is important to note, however, that DMFCs release carbon dioxide during their operation in the same way as a fossil fuel would.

FUEL CELL COSTS

The cost of a fuel cell depends on its application and on the type of fuel cell technology. Fuel cells for automotive applications are expected to be the cheapest available. However, these may have shorter lifetimes than those aimed at large-scale stationary power generation or power and heat generation. Production volumes for the latter are likely to be much smaller too, so the economies of scale achieved from mass production of the former may not be available.

The cheapest fuel cells at the beginning of the second decade of the 21st century are PEM fuel cells for automotive applications, which have a cost of around \$50/kW, as shown in Table 7.2. This is based on a U.S. Department of Energy estimate for an 80 kW unit based on a volume production of around half a million each year.

For stationary applications including domestic heat and power generation and larger-scale power generation the system demands are more stringent and costs higher. Fuel cell stacks in the capacity range of 1–10 kW for domestic systems built from either PEM fuel cells or SOFCs cost around \$750/kW, but these units require significant additional equipment and this pushes the unit cost for a complete installation to \$1500–2000/kW.

Larger stationary applications, generally for both heat and power, can be met by large PAFC installations or large MCFC installations. For an installation around 1 MW, the cost is likely to be around \$4000/kW. SOFC units can also be used for large stationary applications but current costs appear to be higher at \$7000/kW. This is similar to the cost of a MCFC designed to burn biogas.

TABLE 7.2 Fuel Cell Costs

Fuel Cell Type	Cost (\$/kW)
Alkaline fuel cell	9000
Phosphoric acid fuel cell	4000
PEM fuel cell for automotive applications	50
PEM fuel cell for domestic applications	1500–2000
Molten carbonate fuel cell burning natural gas	4000
Molten carbonate fuel cell burning biogas	7000
Solid oxide fuel cell for domestic applications	1500–2000
Solid oxide fuel cell for stationary applications	7000
Direct methanol fuel cell	Typically around 50,000 for small portable generator

Alkaline fuel cells have generally been considered too expensive for all but niche applications. Representative costs are difficult to find for these units. The most cost effective appears to have a cost around $9000/kW. Most expensive of all, however, are DMFCs. Small units with capacities up to 100 W can be bought for around $5000, or roughly $50,000/kW.

Chapter 8

Hydropower

Hydropower is probably the oldest renewable energy resource in the world and one of the first sources of mechanical power. The earliest known reference is found in a Greek poem from 85 BCE and there are references in Roman texts too. Simple wheels used to drive mills and grind grain were known in China during the 1st century, and by the beginning of the second millennium the technology was widely known throughout Asia and Europe.

Early mills used wooden paddles but iron was introduced in the 18th century during the Industrial Revolution in England. Innovation soon led to the development during the 19th century of many of the turbines now in use in modern hydropower stations.

Hydropower capacity grew strongly during the 20th century and until late in that century it was the only significant renewable source of electrical power. According to the Renewable Energy Policy Network for the 21st Century (REN21) Global Status Report for 2013[1] total global hydropower capacity at the end of 2012 was 990 GW excluding pumped storage hydropower capacity. This is an increase of 115 GW compared to the estimate of global capacity from the World Energy Council for the end of 2008 of 875 GW, as shown in Table 8.1. According to the International Hydropower Association, global capacity includes at least 11,000 power stations and 27,000 generating units. REN21 put total global electricity generation from hydropower in 2012 at 3700 TWh, around 15.5% of total global electricity generation, which stood at 22,500 TWh in 2012 according to the BP Statistical Review of World Energy.

Hydropower is widely distributed and few regions are without significant hydropower potential. The countries of the developed world have exploited many of their best sites already and hydropower generation forms part of the bedrock upon which developed nations' prosperity is based. Development elsewhere has been slower but there have been major advances in Asia, particularly in China, in recent years, and many of the countries of South America rely heavily on hydropower for electricity generation. Even so, most of these regions have much remaining capacity, while in Africa hydropower is significantly underdeveloped. Table 8.1 shows the breakdown of global capacity at the

1. Renewables 2013 Global Status Report, REN21, 2013.

TABLE 8.1 Regional Installed Hydropower Capacity

	Installed Hydropower Capacity (GW)	Percentage of Global Total
Asia	307	35%
Europe	221	25%
North America	168	19%
South America	132	15%
Africa	22	3%
Oceania	14	2%
Middle East	11	1%
Total	875	100%

Source: World Energy Council.

end of 2008 based on figures from the World Energy Council's 2010 Survey of Energy Resources.

In spite of its potential and obvious advantages, the development of hydropower can often be difficult, particularly where large projects are concerned. Major hydropower projects are often extremely disruptive and, if not developed sensitively, they can lead to a range of environmental problems. Large hydropower plants, particularly those involving dams and reservoirs, will inevitably change the environment in which they are constructed, leading to displacement of people and wildlife and the destruction of ecologies. With care these changes can be managed, but careless and sometimes reckless development during the second half of the 20th century led to hydropower acquiring a bad reputation during the 1980s and 1990s.

Since then the industry has made an effort to reform its practices, and the World Commission on Dams addressed the main problems in *Dams and Development: A New Framework for Decision Making*.[2] This report proposed a complete reassessment of the criteria and methods used to determine whether a large hydropower project should be constructed. It also laid out an approach to decision making that took account of all the environmental and human rights issues that a project might raise—an approach that should potentially filter out bad projects but allow well-conceived projects to proceed.

2. Dams and Development: A New Framework for Decision Making, World Commission on Dams, Earthscan, 2000.

When projects are well designed and construction is carried out carefully, large hydropower schemes have the potential to transform the lives of those who benefit from them. Many such schemes provide water for irrigation and drinking as well as power, and they can allow new industries to be established too.

Economically hydropower is considered expensive to build but, when accounted for correctly, it can become one of the cheapest sources of electricity available. Since 2000, the introduction of large quantities of renewable generation from wind and solar power have also led to the recognition that hydropower has an important role to play in the balancing of intermittent renewable generation on grid systems. This is leading to a further reassessment of the role of hydropower. Pumped storage hydropower plants, which are large energy storage plants based on hydrotechnology, can be used to store energy from renewable plants for use when needed. However, conventional hydropower can provide significant grid support for other renewable generation too.

Large hydropower projects—those over 30 MW is size—are not generally considered by regulatory authorities to be new renewable generation and in most regions do not attract support such as grants, special tariffs, or tax breaks. However, smaller hydropower schemes, which are generally classified as "small hydropower," will often be included among the technologies that attract such support mechanisms. These smaller schemes are also less disruptive than their larger relatives and are consequently much easier to build.

HYDROPOWER RESOURCE

The energy that is extracted from water by a hydropower plant and converted into electricity is potential energy contained within the mass of water as a consequence of its elevation. This energy is released as the water flows downhill, normally being dissipated in various ways within the watercourse down which it flows. A hydro-turbine can extract some of this energy and use it to produce electric power.

The water flow is created, in the final instance at least, by rainfall, but the water vapor in clouds that is the source of rain is raised into the atmosphere by the effect of the sun's heat on the world's seas and other water surfaces. Solar energy generates the heat that vaporizes the water and lifts it into the high atmosphere, generating the potential energy that is available later for release. Therefore, hydropower is ultimately a form of solar power.

When estimating how much electricity might potentially be available from hydropower resources, a number of measures are commonly used. One is the gross theoretical hydropower potential. The gross theoretical hydropower potential of a region is the total amount of energy that would be released each year if all the energy contained in rain falling across the region was exploited to sea level at the borders of the region. Carrying out the calculation to sea level maximizes the energy available.

TABLE 8.2 Regional Hydropower Potential

	Gross Theoretical Hydropower Capacity (TWh/year)	Technically Exploitable Hydropower Capacity (TWh/year)
Asia	16,618	5590
Europe	4919	2762
North America	5511	2416
South America	7541	2843
Africa	3909	1834
Oceania	654	233
Middle East	690	277
World total	39,842	15,955

Source: World Energy Council.

Table 8.2 contains figures for the gross theoretical hydropower capacity for all the major regions of the world. Asia has the greatest potential at 16,618 TWh/year, followed by South America (7541 TWh/year), North America (5511 TWh/year), Europe (4419 TWh/year), and Africa (3909 TWh/year). Potential in the Middle East (690 TWh/year) and Oceania (654 TWh/year) is much more limited.

There are technical, economic, and environmental reasons why this gross capacity can never be fully realized. A second measure, the technically exploitable hydropower capacity, provides a more realistic figure for the amount that might eventually be used. This is a measure of the capacity that could be exploited using currently available technology. Regional technically exploitable hydropower capacities are also shown in Table 8.2. These are significantly smaller than the gross theoretical capacities. Across Asia, the technically exploitable capacity is 5590 TWh/year, 34% of the gross theoretical capacity. Technical capacities in other regions are similarly reduced compared to the gross capacity.

Cost may further reduce the potential capacity since some technically exploitable hydropower sites will be too costly to develop. A further measure of potential capacity that takes account of this is the economically exploitable hydropower capacity. The global economically exploitable hydropower capacity has been estimated to be about 13,100 TWh/year.[3] There may be still further limits on development as a result of environmental concerns or other

3. This estimate is from Eurelectric.

restrictions. One final measure of hydropower potential—the exploitable hydropower capacity—reflects this. The global exploitable hydropower potential is around 10,480 TWh/year.

As noted before, global hydropower generation in 2011 was 3400 TWh, just under one-third of the exploitable potential. On this basis, two-thirds of exploitable global potential remain to be developed. However, while such estimates are useful guides, all these potential figures should be treated as approximations because there is no general consensus about how such estimates should be made.

Regional levels of hydropower exploitation vary widely. The most highly developed region is Europe where, based on World Energy Council figures from 2010, 26% of the technically exploitable potential has been developed, followed by North and Central America with 22%. South America has exploited 20% of its technical potential and Oceania 21%, but Asia has only exploited 13%, the Middle East 10%, and Africa 5%.

Global installed hydropower capacity was estimated to be just under 1000 GW at the end of 2012. The most recent breakdown of this capacity by region is from 2008, and shown in Table 8.1. This shows that Asia has the greatest installed capacity, 307 GW in total. Over half of this (171 GW) was in China. Europe has the second largest gross capacity with 221 GW. The capacity in North America at the end of 2008 was 168 GW, with 73 GW of this in Canada and 77 GW in the United States. South America had 132 GW, dominated by Brazil with 76 GW. Against these figures the total in Africa, 22 GW, reflects the low level of exploitation on that continent. There are also small capacities in both Oceania and the Middle East as shown in the table.

HYDROPOWER SITES

The first stage in building a hydropower plant is to find a suitable site. This may appear obvious, but it is important to realize that hydropower is extremely site specific. Not only does it depend on a suitable site being available, but the nature of the project will depend on the topography of the site. You cannot have a hydropower plant without a suitable place to construct it. In the case of large hydropower projects (> 10 MW in capacity), sites will often be a long way from the place where the power is to be used, necessitating a major transmission project too.

A successful hydropower project requires a river with suitable hydrological conditions. The amount of energy that can be taken from the river will depend on two factors: the volume of water flowing along it and the drop in riverbed level (normally known as the head of water) that can be exploited. The available power increases with the volume of water while a steep riverbed carrying a fast-flowing river will generally yield more electricity than a slowly descending, sluggish one of similar size. For a given volume of water, the energy available will depend directly on the head height, or drop in water level, that can be utilized and this is normally larger the steeper the riverbed.

This does not mean that slow-flowing rivers are not suitable for hydropower development. They often provide sites that are cheap and easy to exploit. In contrast, steeply flowing rivers are often in inaccessible regions where exploitation is difficult.

Hydropower sites vary in potential from a few kilowatts to many hundreds of megawatts. Occasionally sites will yield thousands of megawatts. The largest single developed site in the world is the Three Gorges Dam on the Yangtze River in China, with a generating capacity of 22,500 MW. Probably the largest unexploited site is on the Congo River in Africa where a multiple barrage development is estimated to be capable of supporting up to 35,000 MW of generating capacity. This is exceptionally large; most are smaller. Large projects of this type, where developed, are likely to be multipurpose projects involving flood control, irrigation, fisheries, and recreational usage, as well as electricity generation. Smaller projects may be multipurpose or they may simply generate electricity.

In choosing a site, hydrology is important, but so too are geography and geology. Given a river capable of supplying energy, the optimum site or sites for extracting this energy will be determined by the geography. Once a site has been identified, an extensive geological survey will then be necessary to determine the underlying structures. Many hydropower plants are physically massive and can generate enormous pressures, leading to stresses in underlying strata and potential fractures. These can be disruptive if possible faults are not identified before work begins. Where large reservoirs are involved, more stress can be created, and this has in some cases led to the generation of seismic tremors as underlying strata react.

How does one set about locating a hydropower site? Many countries have carried out at least cursory surveys of the hydropower potential within their territory and provisional details of suitable sites are available from the water or power ministries. Sometimes much more detailed information is available but this cannot replace an onsite survey. Indeed surveys carried out as part of a feasibility study form an integral of any hydropower scheme. For a large scheme a feasibility study may account for 1% or 2% of the total cost. For smaller schemes it can reach 50%.

CATEGORIES OF HYDROPOWER PLANT

Hydropower plants are traditionally broken down into categories depending on their size. The usual categorization is shown in Table 8.3. The smallest plants, with capacities between 1 kW and 100 kW, are called micro-hydropower plants. Between 100 kW and 1 MW a plant is described as a mini-hydropower plant. Small hydropower plants are generally those with capacities between 1 MW and 10 MW, but this upper limit can vary from country to country and in some cases may be as high as 30 MW. Plants with capacities larger than

TABLE 8.3 Hydropower Plant Categories

Micro	1–100 kW
Mini	100 kW–1 MW
Small	1 MW to 10–30 MW
Large	Above 10–30 MW

Source: Private mini-hydropower development study: The Case of Ecuador, UNDP/World Bank, 1992.

10 MW (or up to 30 MW depending on jurisdiction) are classed as large hydropower plants.

Sometimes an intermediate category for Medium hydropower plants is also introduced between small and large hydropower. If used, this is typically for plants between 5 MW and 50 MW; those above are large and those below are small. From a global perspective large hydropower is the most important category and accounts for most of the hydropower capacity in operation today.

Technically, large hydropower plants are the most sophisticated and are generally individually designed for each site using turbines that have also been made specifically for the power plant. Small hydropower plants are similar to large plants but some use off-the-shelf turbines and other components rather than bespoke components. Mini- and micro-hydropower installations usually employ standard turbines and many involve novel, often cost-effective, designs not used in larger plants.

LARGE HYDROPOWER PLANTS: DAMS AND BARRAGES

Once a potential site for a hydropower scheme has been identified, there are two common ways of exploiting it. The first, called a run-of-river scheme, does without a reservoir, though it will usually involve some sort or barrage across the waterway. Instead it takes water directly from the river though channels and pipes it to the power house where the turbines are installed. The second is to build a dam and create a reservoir behind it from which water is taken to drive one or more hydraulic turbines that are installed in the project's power house. This can be situated at the base of the dam structure but may also be in places some distance downstream of the dam to exploit the maximum head of water possible.

Run-of-River Scheme

A run-of-river scheme is the simplest and cheapest hydropower project to develop. Since it requires no dam, a major constructional cost is avoided. Geological problems associated with dam construction (see next section) are

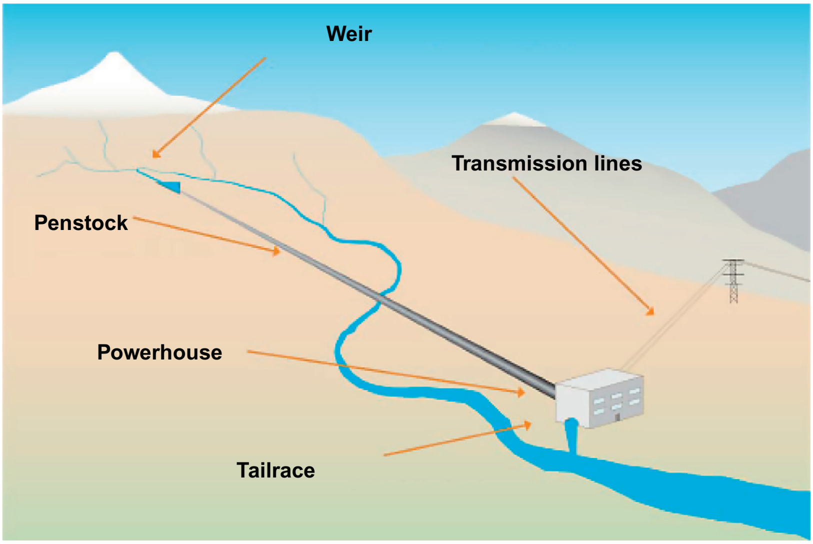

FIGURE 8.1 Layout for a small run-of-river hydropower project.

avoided too. However, some sort of diversion structure will be required to direct water from the river into a canal called the headrace, which carries the water to a spot above the power house of the plant, exploiting the local topology to achieve the minimum drop in elevation (Figure 8.1). From here the water falls steeply through a pipe called the penstock into the power house and the turbines. (In smaller projects the headrace and penstock may be the same pipe.) The height of the penstock intake above the power house represents the head of water that is available for generating power in the plant. It is the pressure at the bottom of this head of water, dependent on the height of that head, that provides the force to drive the turbines. Having passed through the turbines the water is returned to the river at a point downstream of the diversion structure through another pipe called the tailrace. For large run-of-river plants the headrace can be tens of kilometers long.

The simplicity of the run-of-river scheme is attractive but it is also the main weakness of this type of development. With no dam to conserve water, the power plant must rely exclusively on the flow of water in the river. As this fluctuates, so will the amount of power that can be generated. Under drought conditions the plant will be able to generate no power, whereas when the river is in flood, much of the available water will have to be allowed to flow past the diversion system without being exploited. The same applies when power from the plant is not needed for the grid. Nevertheless, this type of project does have significant advantages besides cost, particularly because of the small amount of environmental disruption it causes.

Run-of-river hydropower plants are typically between 10 MW and 1000 MW in generating capacity. They could be larger, in theory, but in practice larger plants of this type have not been built. A series of run-of-river power plants along the same river can exceed 1000 MW in capacity.

Dam and Reservoir Projects

The alternative to the run-of-river is a dam and reservoir project. This will involve a major civil engineering undertaking: the construction of a dam.

The purpose of a dam is to create a reservoir of water that builds up behind it. The reservoir is essentially a form of energy storage system. Once created, the reservoir allows some measure of control over the flow of water in the river beyond the dam and consequently the flow through the turbines in the power house. Water can be conserved during periods of high flow and used up when rainfall is low. A dam can also be used for flood control. More recently, dam and reservoir hydropower systems have been used to help balance intermittent forms of renewable energy on the electricity grid.

If a dam is to be constructed, then a very careful geological survey of the underlying rock will be needed to identify any faults that might make it unstable or allow water to flow beneath it. Geological faults or unsuitable substrata need not prevent construction of a dam because they can be treated, but if they are only discovered during construction, or later, they are likely to result in massive additional costs and delays.

The layout of a dam and reservoir scheme is essentially the same as for a run-of-river plant. Water is extracted from the dam and carried though a headrace until it is above the power house where it enters the penstock and falls into the turbines. From the turbine hall it is then returned to the river through a tailrace. However, since the dam itself may generate the main head of water, the design can often be more compact than for a run-of-river plant and a headrace dispensed with Figure 8.2.

The construction of a dam is a complex engineering project. Water flowing down the river must be temporarily diverted or coffer dams must be erected to isolate part of the riverbed so that construction can take place. The size and complexity of dam construction means that this part of the project will account for up to two-thirds of total project costs.

While the dam forms the major part of the construction project, the reservoir behind the dam is likely to have the largest environmental impact. The lake created behind the Three Gorges Dam is 600 km long, but even small dams can create large areas of water with dramatic effects on the local environment.

DAM TYPES

There are three principle types of dam used for hydropower projects: embankment dams, concrete gravity dams, and concrete arch dams. The simplest of these and the cheapest to construct is the embankment dam.

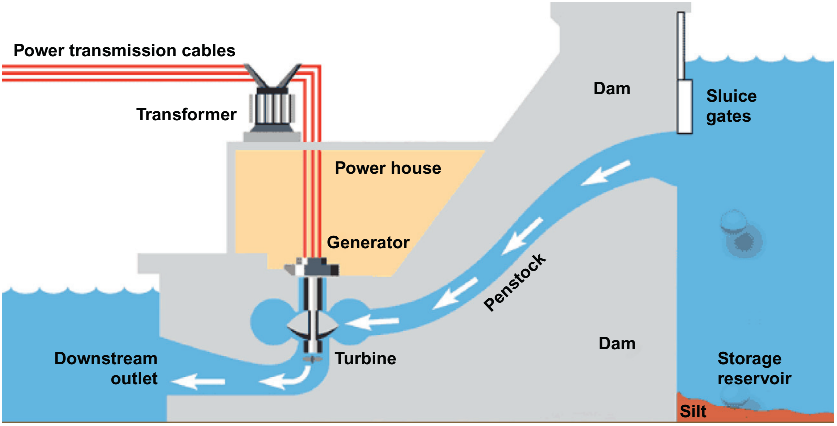

FIGURE 8.2 Dam and reservoir power hydropower development with power house integrated into dam structure.

Embankment Dam

Embankment dams come in two principle varieties: earthfill embankment dams and rockfill embankment dams. An earthfill embankment dam is made by building a foundation wall that is embedded into the rock below the dam to prevent water flowing beneath it and then creating a core of impermeable clay on top of this (Figure 8.3). Above this the remaining structure is built from earth, normally from the surrounding area. It is the mass of earth and clay that holds the dam in place. A rockfill dam uses rock instead of earth, normally with an impermeable layer on the upstream face of the dam to prevent seepage through the porous core. However, both types of dam can tolerate controlled amounts of seepage.

Both types of embankment dam are vulnerable to overflow of water eroding the dam structure, so each must be provided with a spillway that can release

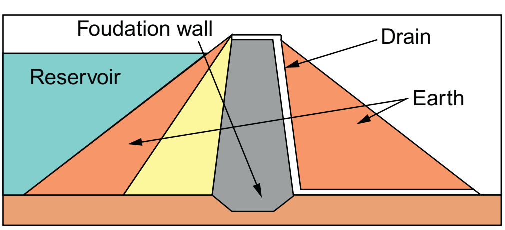

FIGURE 8.3 Cross-section of an earthfill embankment dam.

water from the reservoir behind it if the water level becomes too high. These dams have shallow sloping sides to create a large mass and are usually employed in shallowly sloping valleys, creating relatively wide and shallow reservoirs.

Concrete Gravity Dam

The second main type of dam is the concrete gravity dam. Like the embankment dam this relies on its mass to resist the pressure of water, but resistance may also be aided by the concrete structure being bedded into the underlying rock and, where possible, into the sides of the watercourse where it is constructed (Figure 8.4). Concrete dams cannot deform like embankment dams and are vulnerable to water passing beneath them, which may undermine the structure, so the integrity of the geological strata below the dam is essential.

A variant of the concrete gravity dam is the concrete buttress dam. This has a vertical upstream face but downstream it is supported by massive triangular buttresses built into the rock below. Such dams require less concrete than a concrete gravity dam since they use the buttresses to resist the water pressure.

Concrete Arch Dam

The final type of dam is the concrete arch dam. This deploys the principle of the arch to resist the pressure of water. An arch dam can only be built in a steeply sided rock ravine where both the sides of the ravine and the underlying substrata are sound. The arch is built on its back, bowing upstream, with the sides bedded into the rock faces of the side of the ravine (Figure 8.5). The concrete arch dam

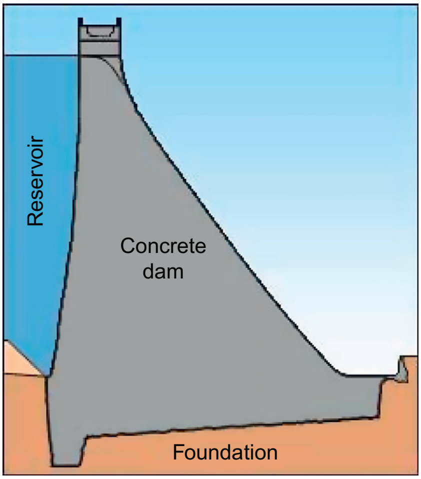

FIGURE 8.4 Cross-section of a concrete gravity dam.

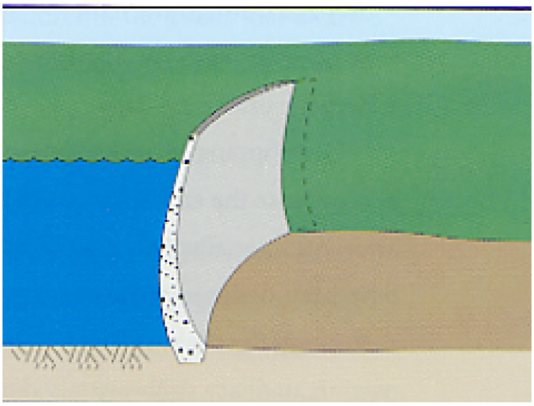

FIGURE 8.5 Cross-section of an arch dam.

is the most precisely engineered of all dam types and requires less material than other types and where it can be built it is the strongest of dams.

HYDROPOWER TURBINES

The most important mechanical component of any hydropower-generating unit is its turbine. This is the device that converts the energy contained in the moving water into the rotary motion necessary to turn a generator and produce electricity. The hydraulic turbine is a simple, reliable, and well-understood component made from simple materials. Most turbines are made from iron or steel or cast from bronze or other alloys. In the past wood was commonly used too.

The history of the hydropower turbine is long. The earliest known, water wheels for grinding grain, were used by the Romans and were known in China in the 1st century. They were common across Europe by the 3rd century and could be found in Japan by the 7th century. The Doomsday Book of AD1086 records 5000 in use in the south of England. These early water wheels were made of wood. Iron was first used in the 18th century by an English engineer, John Smeaton.

Early water wheels were simple paddle wheels, the lower edge of which was placed into a flowing stream where the pressure of water against one side of the immersed paddle caused it to turn. This is the basis for one group of turbines in use today called reaction turbines. It was later discovered that more energy could be extracted by damming the water to create a head and using the pressure at the bottom of this head to create a jet of water that was directed against the paddles of the wheel. This principle led to the development of the second main branch of hydropower turbines, the impulse turbines.

Modern turbines can be extremely efficient with the best converting 95% of the energy contained in the water into rotary motion. They can be stopped and started very quickly and are capable of operating reliably for many years. Large

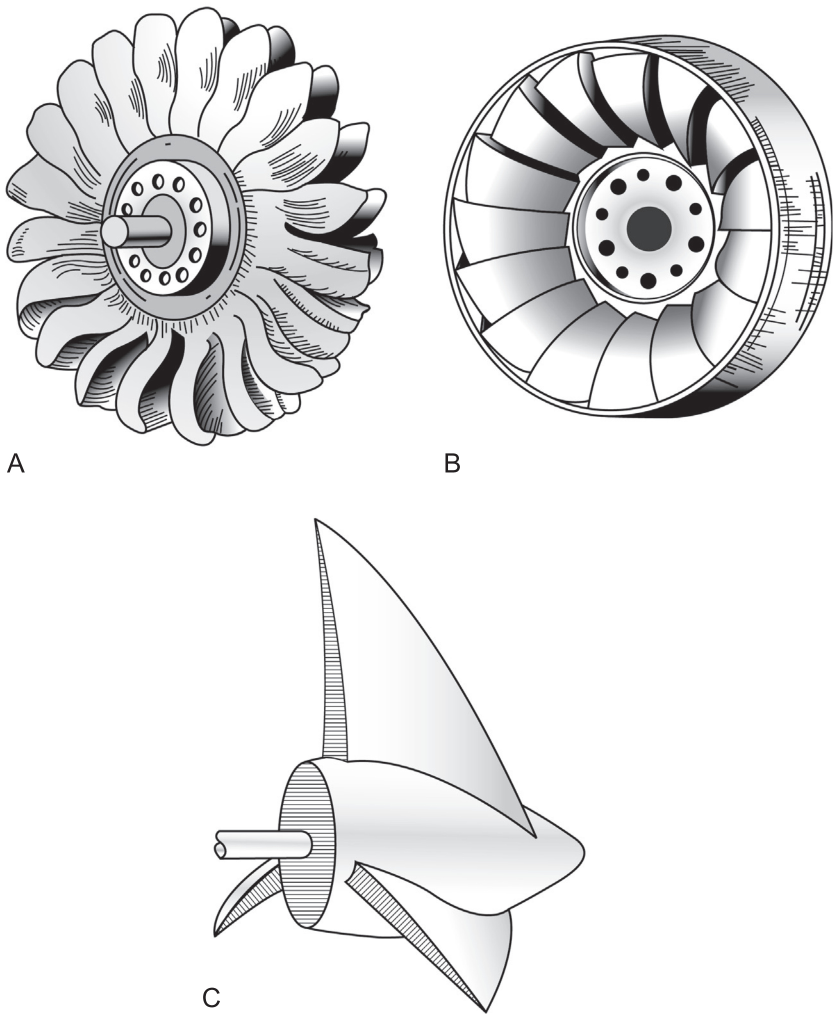

FIGURE 8.6 Hydropower turbines: (A) Pelton (B) Francis, and (C) propeller turbines.

hydropower projects generally make use of one of a small number of turbine types, discussed in the following sections (Figure 8.6). Small hydropower plants utilize these same types but they can exploit a wider range of turbines including some extremely novel designs (see later).

Impulse Turbines

A tall column of water creates a high pressure at its base and if the water under high pressure is released through a fine nozzle it will create a high-speed jet of water. If this jet is directed onto a bucket-shaped paddle on a wheel it will

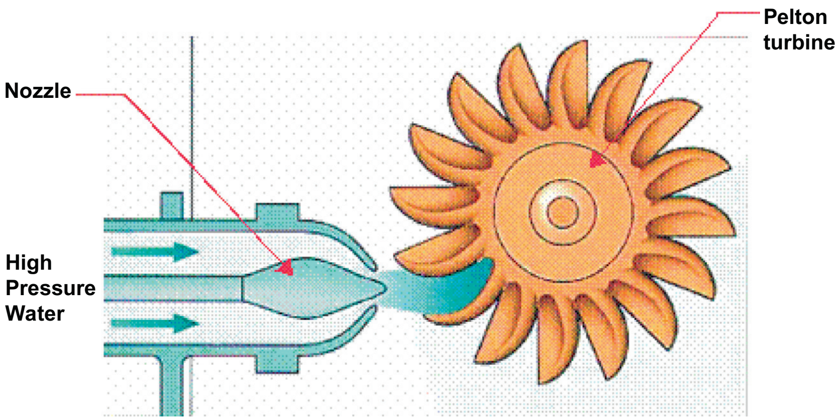

FIGURE 8.7 A Pelton turbine and nozzle.

generate an impulse that causes the wheel to turn. This is the basis for the impulse turbine.

The main type of impulse turbine in use today is the Pelton turbine, patented by the American engineer Lester Allen Pelton in 1889 (Figure 8.7). The Pelton turbine operates at its highest efficiency, 95%, when the speed of movement of the bucket on the wheel rim is half that of the water in the jet directed into it.

The Pelton turbine is generally used where a high head of water is available and the flow rate is low. For large hydropower plants they are normally preferred when a head height of more than 450 m is available, but they can be used for heads a low as 200 m. The maximum head for a single Pelton turbine is normally 1000 m. Beyond that the fall must be divided into two sections and more than one turbine will be required.

A simple Pelton turbine will have one nozzle but power output can be increased by using up to four nozzles directed at the same wheel. Most Pelton turbines are mounted vertically with a horizontal axis of rotation, but they can be mounted with a vertical axis so that the turbine lies horizontally. A key feature of all Pelton turbines is that they must operate in free air, not submerged. The turbine must, therefore, always be positioned above the water level at the bottom of the head of water being utilized for power generation.

The speed at which a Pelton turbine rotates will be determined by both the flow rate of water directed into its buckets and the load into which it is feeding. If the load falls, the turbine will speed up. This can be controlled by reducing flow through the nozzles. The optimum efficiency is achieved when the turbine is operating between 60% and 80% of maximum load.

A second type of impulse turbine available for high-head use is the Turgo turbine. This is similar to the Pelton in design, but whereas with the Pelton turbine the water jets are in the same plane as the turbine wheel, in the Turgo turbine the jet strikes each bucket from one side and then exits the turbine at the other side. The Turgo turbine can handle higher flow rates than the Pelton but is generally more difficult to construct. It is normally used for medium-head applications between those best suited to the Pelton and those more suited to reaction turbines.

Reaction Turbines

For heads of water below 450 m, a reaction turbine will normally be the first choice. This type of turbine must be completely submerged to operate efficiently. Whereas the impulse turbine harnesses the kinetic energy (and the momentum carried by its mass) of a jet of high-pressure water, a reaction turbine responds to the pressure (potential energy) created by the weight of water at the base of the head acting on one side of its blades.

There are several different types of reaction turbine. The most popular, accounting for 80% of all hydraulic turbines in operation, is the Francis turbine. This can be used in almost every situation, but for very low heads, propeller turbines and Kaplan turbines are frequently preferred.

Francis Turbine

The Francis turbine was developed by James Bichens Francis around 1855. Its key characteristic is the fact that water changes direction as it passes through the turbine. The flow enters the turbine in a radial direction, flowing toward its axis, but after striking and interacting with the turbine blades it exits along the direction of that axis (Figure 8.8). It is for this reason that the Francis turbine is sometimes called a mixed-flow turbine. For it to operate efficiently, water must reach

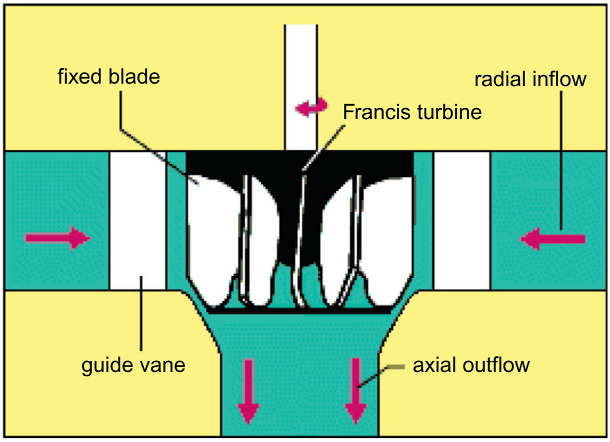

FIGURE 8.8 Cross-section of a Francis turbine.

all blades equally and flow is controlled by a set of valves or gates that curl around the turbine itself in a spiral shape.

The blades of a Francis turbine are carefully shaped to extract the maximum amount of energy from the water flowing through it. Water should flow smoothly through the turbine for best efficiency. The force exerted by the water on the blades causes the turbine to spin and the rotation is converted into electricity by a generator. Blade shape is determined by the height of the water head available and the flow volume. Each turbine is designed for a specific set of conditions experienced at a particular site. When well designed, a Francis turbine can capture 90–95% of the energy in the water.

The Francis design has been used with head heights from 3 m to 600 m, but it delivers its best performance between 100 m and 300 m. Flow rate is often the limiting factor for a given head. As the head height rises, the size of the turbine must fall, making fabrication more difficult. High-head Francis applications, therefore, require a large flow to be successful. Conversely, for low-head applications the flow must be low or the turbine will become excessively large. It is for this reason that while the Francis turbine is the most versatile, different designs are generally used for both very high and very low heads.

Francis turbines are also the heavyweights of the turbine world. The largest, found at both the Itaipu power plant on the Brazil–Paraquay border and at the Three Gorges Dam in China, have generating capacities of 700 MW each.

Propeller and Kaplan Turbines

If the head of water becomes too low, the rotational speed of a Francis turbine falls and with it efficiency. For low-head applications an alternative turbine type is required and the most successful is the propeller turbine (Figure 8.9).

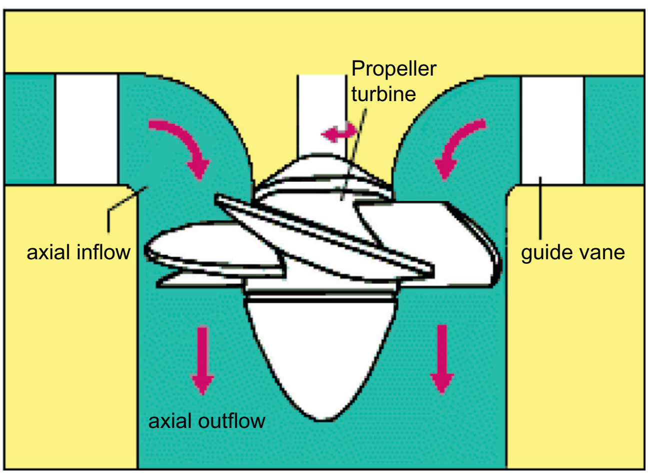

FIGURE 8.9 Cross-section of a propeller turbine.

A propeller turbine looks like the screw of a ship, but its mode of operation is the reverse of the ship's propulsion unit. In a ship a motor turns the propeller, which pushes against the water, forcing the ship to move. In the hydropower plant, by contrast, moving water drives the propeller turbine to generate power.

Propeller turbines are most useful for low-head applications such as slow-running, lowland rivers. The head of water is typically 10 m or less. Their efficiency drops off rapidly when the water flow drops below 75% of the design rating, so plant designers often use multiple propeller turbines in parallel, shutting down some when the water flow drops to keep the remaining turbines operating at their optimum efficiency.

A variant of the propeller turbine is the bulb turbine, used for extremely low-head conditions. In this design the turbine is integrated with a water-tight generator and enclosed in a bulb-shaped container. The turbine rotor can have fixed or variable blades. Water flows into one end of the bulb-shaped container, called the nacelle, and out the other, with no change of direction. The use of the nacelle helps concentrate the flow to maximize energy capture. The bulb turbine has been used in tidal power plants.

There are cases where multiple turbines cannot easily be used. In these cases, an alternative called the Kaplan turbine has sometimes been utilized instead. This low-head variant of the propeller turbine was developed by Viktor Kaplan in 1915. Its primary feature is a set of blades that can be adjusted to maximize efficiency under different flow conditions. These turbines can operate at higher flow levels than conventional propeller turbines and are suited to heads between 10 m and 50 m. A variation of the Kaplan turbine is the diagonal-flow turbine that can operate at higher rotational speeds than the Kaplan and is suited to higher heads.

Deriaz Turbine

A further turbine available to hydropower plant designers is the Deriaz turbine invented by Paul Deriaz. This has blade shapes similar to a Francis turbine but these blades are adjustable and so they can be adapted to different flow rates. The Deriaz turbine is another mixed-flow turbine in which water enters from the size but exits along the turbine axis. It is best suited to head heights between 20 m and 100 m, in between the ranges of Francis and Kaplan turbines.

GENERATORS

The turbine in a hydropower plant is connected directly to a generator to produce electricity, as shown in Figure 8.10. Generators for large hydropower turbines are normally synchronized to a grid at 50 Hz or 60 Hz and this controls the speed at which both the turbine and generator must rotate.

Turbine speed will vary with turbine type. For a Pelton turbine the typical rotational speed is 400–1000 rpm. Francis turbines typically rotate between

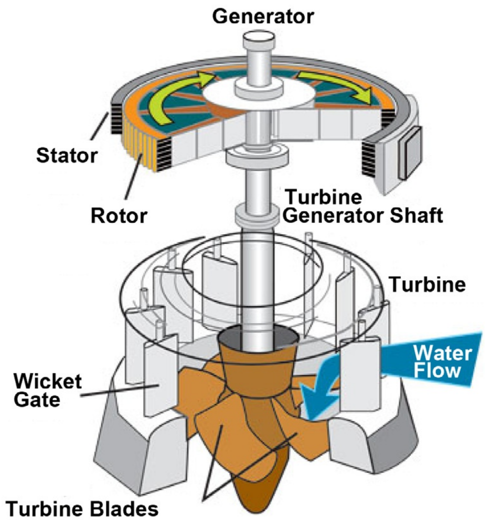

FIGURE 8.10 Propeller turbine and generator.

100 rpm and 1500 rpm. Propeller turbines, meanwhile, run at speeds between 60 rpm and 300 rpm. In each case the generator must be tailored to suit the turbine speed.

A synchronous generator cannot operate at any speed but instead must turn at one of a number of fixed speeds. Flexibility of generator rotational speed is achieved by varying the number of poles. The smaller the number of poles, the higher the rotational speed needed to achieve a given synchronization frequency. For 50 Hz operation a two-pole generator must rotate at 3000 rpm, while a 60 Hz machine must rotate at 3600 rpm. For eight-pole generators the necessary speeds would be 750 rpm and 900 rpm. Therefore, generators for slow-speed turbines require a larger number of poles.

While most large hydropower plants rely on synchronous generators of this type, an increasing number of smaller units are starting to employ variable-speed generators. These can operate at any speed, the latter being determined to optimize efficiency. Output from the generator is then converted to grid frequency using solid-state power electronic devices. The power-handling capabilities of the latter generally limit the generator size to tens of megawatts.

SMALL HYDROPOWER

Small hydropower plants are defined as much by regulatory regimes as by their inherent design features. Plants below a certain size usually qualify as new renewable generation and these are called small hydropower plants. The limit varies from country to country. In Sweden, for example, the upper size of a small hydropower plant that can attract support is 1.5 MW. In Italy the limit is 3 MW, in France 12 MW, in the United Kingdom 20 MW, and in Canada and the United States it is between 30 MW and 50 MW.

The global installed capacity of small hydropower plants is not known precisely but is probably around 50,000 MW, although it could be significantly higher. The potential for future development is equally uncertain but the International Energy Agency has suggested that there may be 150,000–200,000 MW capable of development.

The design of a small hydropower plant depends very much on its size. Those in the small hydropower category of Table 8.3 (1 MW to between 10 MW and 30 MW) will be approached in a similar way to a large hydropower project. At this size, dam construction is less likely to be cost effective but some sort of barrage may be employed. At the smaller end of the range off-the-shelf rather than site-specific turbines are also likely to be used. In most cases these will be the types discussed earlier for large hydropower schemes.

Mini- and micro-hydropower plants of less than 1 MW in size are likely to be approached differently. Here cost becomes the overriding concern and a range of novel techniques including the use of cheap pumps as turbines and inflatable barrages may be employed to keep costs down.

One major difference between large and small hydropower is the breakdown of head height into categories. For a small hydropower plant a head above 100 m will be considered a high head and any project with a head of this or higher will employ a high-head turbine such as a Pelton turbine. For very small projects a Pelton turbine may be used at even lower head heights. Projects with heads between 30 m and 100 m are classified as medium-head schemes, while anything under 30 m qualifies as a low-head plant.

Plant design will be much simpler in a small hydropower scheme. Most will be run-of-river (or run-of-stream) and any intake structures, where used, are likely to be rudimentary to keep costs low. For larger plants a weir may be employed. Others will take water directly without any type of barrage. In many cases the turbine generator will be placed directly into the waterway.

If water is extracted from the river or stream it may be carried some distance through the equivalent of a headrace, but more commonly it will be fed directly into a penstock-type conduit that carries it directly into the turbine. Penstock length can affect project costs significantly so this will be kept as short as possible.

Turbine types for small hydropower schemes will depend on head height; Pelton turbines for high-head, Turgo and Francis for medium-head, and propeller and Kaplan turbines for low-head applications. Other turbines are also commonly used. These include the cross-flow turbine–a low-cost type of impulse turbine–the Archimedes Screw and the Gorlov turbine that is a little like a vertical-axis wind turbine that operates under water. Simple paddle-type water wheels are also common (Figure 8.11). For very small applications cheap pumps can be used in reverse to make turbine generators. These are known

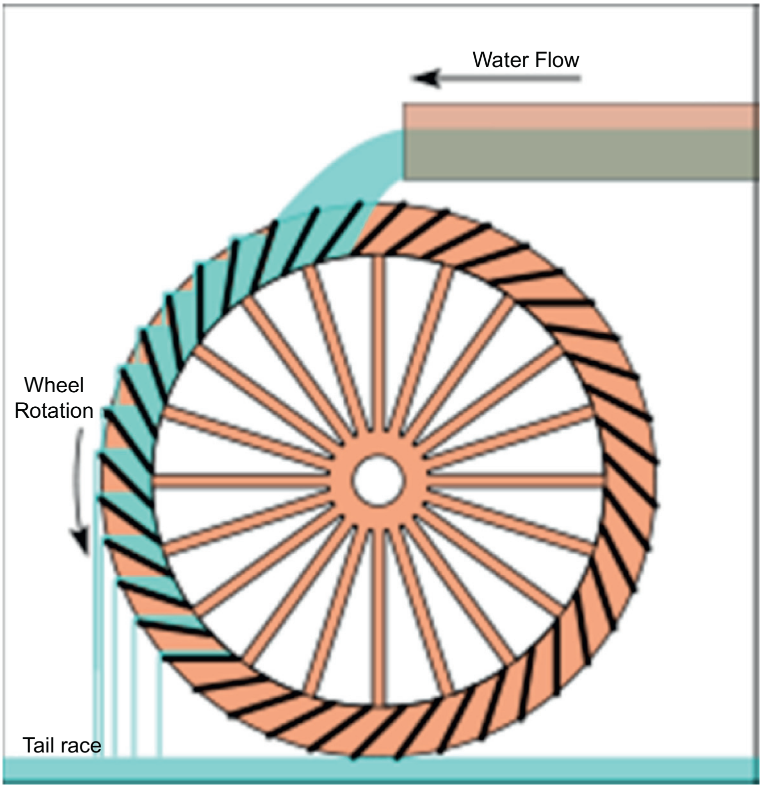

FIGURE 8.11 Cross-section of a water wheel.

as pumps-for-turbines (PATs) and can be used with head heights of 13–75 m to build very cheap hydropower facilities. Small propeller turbines fitted with sealed generators can also be dropped directly into a stream to provide a hydropower-generating system.

While larger small hydropower schemes may use synchronous generators of the sort used in large plants, many small plants employ asynchronous generators that rely on the grid to help them control their speed of operation. In some cases these are simply motors being used in reverse to generate power. The efficiency of such small asynchronous generators is much lower than for large generators.

Small hydropower schemes tend to be relatively more expensive than larger schemes because costs of many of the components do not fall in line with size. The cost of a grid connection can become a large component of the small hydropower scheme and a feasibility study may take a large share of the budget. The extra cost can still be economical if the small hydropower scheme is supplying power directly to consumers where it will be competing with the retail cost of electricity rather than the wholesale cost. Small hydropower schemes can also be extremely effective in supplying power to remote communities far from a grid, especially when the alternative is diesel power.

ENVIRONMENTAL CONSIDERATIONS

The environmental effects of a hydropower project, particularly one involving a dam and reservoir, are significant and must be taken into account when a project is under consideration. What is going to be submerged when a reservoir is created? What effect will the dam or barrage have on sedimentary flow in the river? What are the greenhouse gas implications? Whose interests are affected? For a run-of-river scheme the level of disruption is likely to be lower but extensive environmental studies will still be required. In both cases all these issues must be addressed. Small hydropower schemes are rarely disruptive on the same scale as large hydropower projects and their impact is usually limited so that decisions can often be made at a local rather than a national level. Large schemes have the potential to affect regions and require much more careful scrutiny at regional or national levels.

This problem is not new. Humankind has been altering waterways for at least two millennia and some of the early structures still exist. Roman dams can be found in use in Spain today. In the past dams have been used to provide water for drinking and irrigation and to help control waterways. Only since the end of the 19th century has electricity generation been added to this list of uses.

While dams will always change the environment, in the past the changes wrought have generally been considered positive, providing an improvement in the living standards and conditions of the people affected. Greater environmental awareness coupled with some careless developments led to a change in perceptions toward the end of the 20th century and since then large hydropower schemes have become much more difficult to promote and build. This prompted work by the World Commission on Dams to look at what made a good and what made a bad hydropower project.

The work of the World Commission on Dams resulted in 2000 in the publication of *Dams and Development: A New Framework for Decision Making*.[4] Since then a reappraisal of large hydropower has taken place resulting in greater acceptance of the need to develop waterways. In particular, the link between water projects and the standard of living of people affected has been recognized. This is particularly significant in Africa where hydropower development is far behind other continents and standards of living are often very low too.

So while it is recognized that a hydropower project, particularly a large one, will be disruptive, it is also recognized that it need not necessarily be destructive. Environmental changes will take place, people may be displaced, and

4. Dams and Development: A New Framework for Decision Making, World Commission on Dams, Earthscan, 2000.

habitats destroyed but all these effects *can* be handled sensitively so that, for example, displaced communities are given a stake in the project and much improved living conditions and inundated habitats are recreated alongside the area destroyed.

ENVIRONMENTAL ASSESSMENT

Today, to make a case for a major hydropower project, a thorough environmental assessment will usually be necessary and in most cases it will be mandatory. The effects of the project, including any necessary resettlement, effects on biodiversity, the potential for seismic activity, and the impact on areas downstream of the project, will all have to be evaluated. Such a study should include proposals for the mitigation of any negative effects of the development. In many cases, particularly where international lending agencies are involved, a project will not be permitted to proceed unless the environmental assessment is favorable. This is equally true of public sector and private sector projects.

Resettlement

The most divisive effect of any large hydropower project is likely to be the need to resettle people whose homes or communities will be destroyed. In building the Three Gorges Dam, the Chinese government moved 1.3 million people and more people might need to move if the reservoir banks become unstable. This represents one of the largest resettlement programs for a hydropower project, but even with much smaller numbers the result will be extremely disruptive for those involved.

Resettlement numbers can be large as in the case of the Three Gorges Dam, but how is one to judge if they are too large? One way of at least comparing projects is to determine the number displaced for each megawatt of generating capacity installed. For the Three Gorges Dam this ratio was 71. The Kedung Ombo Dam in Indonesia, a 29 MW project that led to the displacement of 29,000 people, had a ratio of 1000 people/MW. In contrast, the Grand Coulee Dam built in the 1930s in the United States had a ratio of 2.

If people are to be displaced then a rule of thumb for modern developments is that they should be better off economically afterwards than before. More than that, people being moved should have a large say in where they are moved to and a stake in the project. If the project can provide wide community benefits then people will support it. If not, then development should be questioned.

The situation becomes more difficult when whole communities and their cultural and religious sites are likely to be destroyed. It is sometimes possible to move such sites; the most high-profile example of this is the rebuilding of an Egyptian temple before inundation of the Nile behind the High Aswan Dam in Egypt. However, cultural and religious beliefs may make such a solution unacceptable. In many cases, particularly in remote parts of the world, such considerations are all too easily ignored.

Biodiversity

Even if a dam and reservoir does not displace many people it can still have an enormous impact if it affects a large area. The relative impact in this case can be crudely assessed by calculating the area inundated for each megawatt of generating capacity. This ratio for the Three Gorges Dam was 317 ha/MW while for the Grand Coulee Dam it was 5 ha/MW. Meanwhile, the Kompienga Dam in Burkina Faso achieved the score of 1426 ha/MW, the highest of any recent project. Again this is only a broad indication of the effect since it will also depend on the type of terrain that is being submerged. However, large, shallow reservoirs will always have more impact by this measure than deep, narrow ones.

The greatest danger to biodiversity is that a project will destroy the home of an endangered species. Since hydropower projects take a long time to develop, it is possible to create a new habitat to replace the one that is threatened while work continues on dam construction. This can be relatively straightforward for plant species but can be much more difficult for animal species. However, it is feasible. Indeed, some older projects are now introducing managed habitats that were not considered when the plant was initially built.

The effect of dam and reservoir construction on aquatic species is less obvious but can be equally dramatic. In France several rivers no longer support salmon as a consequence of dams, and in China the Yangtze dolphin was declared extinct in 2006, partly as a consequence of hydropower developments along the river that prevented its movement.

The water in deep reservoirs can become deoxygenated, affecting aquatic life that might otherwise live there. On the other hand, the creation of a reservoir can provide new opportunities for fish species and large reservoirs can allow fish farms to develop, creating a new industry that did not previously exist.

Geological Effects

There is a growing body of evidence that the construction of dams and inundation of reservoirs can generate seismic activity in the underlying strata as a result of the pressures generated at the surface. Such effects are normally only found with large dams, over 100 m high. The activity is generally short-lived but in some cases it can persist. And earthquake in Sichuan province in China in 2008 has been linked to a dam. This earthquake caused the loss of 80,000 lives.

Another danger is of landslips in the region around the reservoir. A case of this type in Italy in 1983 caused a reservoir to overtop, leading to the loss of 2600 lives. Such landslips are not only potentially fatal, as in this case, but they also reduce the volume of the reservoir and therefore its utility.

Sedimentation and Downstream Effects

All rivers carry a load of small particles that are borne downstream with the water. When a river is dammed, this load of sediment will often simply settle in the bottom of the reservoir and slowly fill it up. The reservoir for the Sanmen Gorge hydropower plant on the Yellow River in China lost 40% of its volume to sediment in four years. While this is a dramatic case, most reservoirs have to deal with this problem on some scale.

Normally sediment deposition will reach a steady state with enough sediment being carried past the dam to balance that being deposited each year. It is sometimes possible to wash sediment away periodically by opening the sluice gates of the dam. However, sediment is made of abrasive particles and its passage through the plant's turbines will cause wear that may eventually lead to the need for repair or replacement.

Many of the most important environmental effects of sediment deposition are felt downstream of the dam. One immediate effect of loss of sediment is to increase erosion immediately below the dam site. More dramatic is likely to be the effect of loss of sediment on downstream habitats that rely on it. When the Aswan Dam was built on the Nile River it prevented vital sediment reaching the Nile delta. This sediment was the source of the delta's fertility and its loss led both to delta erosion and to a rapid increase in the use of artificial fertilizers. Problems in the Black Sea with algae have also been linked to loss of sedimentary material as a result of dams on rivers such as the Danube.

Against this, a dam on a major river such as the Danube or the Yangtze can help control flooding, making downstream regions much safer. It can also make the river navigable upstream, which can be a benefit for local communities.

Greenhouse Gases

Hydropower projects are generally classified among the power generation schemes with the lowest greenhouse gas emissions. Typical greenhouse gas emissions are 10–13 kg/MWh, similar to that of wind power plants. Not all hydropower schemes are low emitters, however. Some can generate significant quantities of methane, a potent greenhouse gas.

Methane is produced when organic material collects in the bottom of a reservoir where the water is deoxygenated. Under these conditions anaerobic digestion takes place releasing methane gas. To prevent this, project developers should try to remove as much organic material as possible from the region to be inundated before submersion takes place by felling trees and clearing undergrowth where possible. Even so it will be impossible to remove everything.

The Canadian utility Hydro Quebec, which has studied this effect, has found that methane production from reservoirs normally follows a predictable cycle. Production peaks between 3 and 5 years after the reservoir is filled. After 10 years emissions are no greater than for natural lakes. However, there have been cases where much higher levels of methane emissions have been detected.

Interregional Effects

Dam construction can lead to political disputes when rivers cross national boundaries. For example, the damming of the Euphrates River in Turkey has reduced water flow through Syria and Iraq and this has led to frequent disputes. Reduced water flow when water is taken upstream for irrigation or drinking is one problem. Others may relate to sedimentation issues discussed earlier. In all cases, however, friction is likely unless great care is taken with such developments.

HYDROPOWER AND INTERMITTENT RENEWABLE GENERATION

The amount of electricity generation from sources such as wind energy and solar power is increasing in all parts of the world. The output from this type of power generation plant is intermittent, and this causes problems for grid operators who must maintain their grids in balance while absorbing all the energy from these sources when it is available.

The traditional solution to this balancing problem is to maintain fossil fuel–fired plants on standby so that they can be brought into service as the renewable sources fail. Typically some form of gas turbine–based plant will be used to provide this backup. However, where hydropower capacity based on dam and reservoir plants are available, these can often provide both a faster-acting and cheaper means of maintaining the grid in balance.

Pumped storage hydropower plants have provided an energy storage and grid support service for many years (see Chapter 10 on energy storage), but normal hydropower schemes can provide the same service so long as they are operating within their safety limits. A reservoir that is near flood levels will have to lose water whether power is needed or not, and one that is emptied during a dry season may not be able to generate. At other times, however, the plant should be able to stop and start as necessary. Grid support of this type can earn a power plant additional revenue and, as renewable generating capacities increase, hydropower is likely to be seen as increasingly important.

COST OF ELECTRICITY GENERATION FROM HYDROPOWER PLANTS

Hydropower plants are generally considered to be capital-intensive power projects because most of the cost is associated with the construction of the plant and very little with its operation. This means that large amounts of funding must be available at the outset. In the past the high cost of building a hydropower plant was often borne by the public sector, but since the liberalization of electricity markets that started toward the end of the 1980s, it has often fallen to private sector companies to fund them, something that they have often been reluctant to do.

Some major schemes are still funded by the public sector and others are funded through financial-support mechanisms such as the World Bank. Increasingly, too, private investment is finding its way into hydropower. However, financing is often complicated by the fact that at least some of the benefits of a new hydropower project, particularly in the developed world, will accrue to the government. These benefits include flood control and the supply of irrigation and drinking water.

Another factor that affects the economics of hydropower is the long life that can be expected from a well-designed project. While most power plants have useful lives of 30–40 years at most, a hydropower plant can continue to operate for over 100 years provided the turbines are maintained and periodically replaced. However, financing of a project is unlikely to be possible over 100 years, so costs will be weighed heavily on the early years of the project. Once the cost of loans needed to build a hydropower plant have been paid, the cost of electricity from the plant is likely to be as cheap or cheaper than virtually all other sources.

The breakdown of costs for a hydropower plant suggests that typically 60–70% of the total is accounted for by the civil works. Equipment only accounts for 25–35% while engineering and consultancy takes the remaining 5–10%. Since the civil engineering portion of the project will involve considerable labor costs, overall costs will vary with these costs. Labor costs in some regions are likely to be much lower than in others.

Actual capital costs for hydropower plants vary widely but typical costs, based on published figures for recent plants, are between \$1000/kW and \$2000/kW. Many of these plants have been built in developing countries where labor costs tend to be lower than in developed countries. The U.S. Energy Information Administration (EIA) has estimated that the cost of a new 500 MW hydropower plant in the United States, commissioned in 2011 and entering service in 2015, would be \$2134/kW, just outside the upper limit of the preceding range.[5]

Small hydropower plants tend to cost more than the larger projects because many of the costs do not scale with size. Typical costs are from \$1500/kW to as much as \$5000/kW. However, the actual costs of such projects will depend on both size and the type of technology being used. Very small schemes based on pumps as turbines could be a much lower cost.

5. *Assumptions to the Annual Energy Outlook, 2012*, U.S. Energy Information Administration, 2012.

Even with relatively high capital costs, hydropower can offer a low cost of electricity option. For example, in the United States, the EIA estimated that the cost of electricity from a new hydropower plant entering service in 2017 would be $89.9/MWh. Of the common technologies this was only undercut by a natural gas combined cycle power plant without carbon capture and storage. This price will be based on some form of financing and loan repayment. However, for plants in the United States that have paid off their loans, generation costs are estimated to be between $20/MWh and $40/MWh, undercutting virtually any alternative source.[6]

6. See the Pew Center website: http://www.pewclimate.org/technology/factsheet/hydropower.

Chapter 9

Tidal Barrage Power Plants

The tidal rise and fall of the seawater level along a coastline leads to the movement of large volumes of water in and out of coastal inlets and estuaries. This moving water can be used in the same way as the water flowing down a river as a means of generating electrical power.

The simplest way of exploiting the energy available is to build a tidal barrage across the mouth of an estuary or suitable inlet. The tidal changes in sea level will then cause water to flow cyclically backwards and forwards across this barrage. When the tide rises, water flows from the sea into the estuary or inlet, passing through sluice gates in the barrage. At high tide the sluice gates are closed, and when the tide ebbs, the water behind the barrage is allowed to flow back to the sea through hydraulic turbines, generating power in the process (Figure 9.1).

The head height available for generation will vary with the state of the tide and a tidal plant will normally not start generating until sometime after high tide to obtain the optimum head for the site. In principle, it is possible to generate when the tide is rising instead of when it is falling. Usually, however, ebb-tide generation alone is preferred.

Exploitation of tidal motion has a long history and tidal mills with water wheels have been known for the best part of a millennium in Europe and elsewhere. The earliest record is from 900 AD, but there will probably have been much earlier mills in operation. These early mills would impound water during the incoming tide, allowing the mill to operate for about three hours on each tide.

Today's tidal power plants for electricity generation are relatively modern and also rare. Apart from some small plants built in China from the late 1950s onwards, the first commercial plant was built in France where it started operating in 1966. This remains one of the two largest operating tidal power plants and one of only a handful of commercial plants worldwide. The reason why there are so few plants is primarily down to the high cost of building a tidal barrage, the expense of which makes barrage tidal power plants appear uneconomical. There are also a limited number or sites where tidal plants of this type can be constructed, again limiting the potential. Nevertheless, the tidal power plant remains of interest because of its long life and reliability.

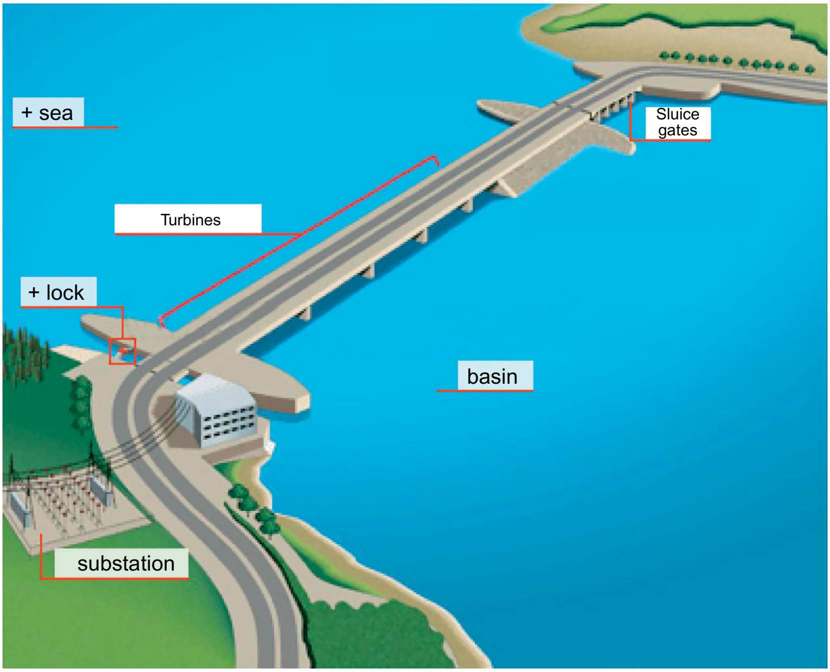

FIGURE 9.1 A tidal barrage power plant.

In addition to the coastal ebb and flow of the tide that is exploited by barrage plants, tidal movement generates coastal currents, which offer another route to exploiting the energy, using tidal current devices that are much like underwater wind turbines. These devices, which have a range of applications, are discussed in Chapter 14 on marine technologies.

TIDAL RESOURCE

The motion of the tides is caused by the gravitational pull of the moon and sun. This motion varies according to a number of cycles. The main cycle is the twice daily rise and fall of the tide as the Earth rotates within the gravitational field of the moon. A second, 14-day cycle is caused by the moon and sun being alternately in conjunction or opposition. This results in spring and neap tides. There are other cycles that add 6-month, 19-year, and 1600-year components but these are much smaller.

Therefore, tidal energy is energy generated through the motion of the planets. As such, it is one of the few renewable energy sources that do not depend either directly or indirectly on radiant solar energy.

The actual size of the tidal movement depends on geographical location. Tidal amplitude in the open ocean is around 1 m but this increases nearer to land. Amplitude can be substantially enhanced by the coastal land mass and by the shape of river estuaries. Under particularly propitious conditions, such as those found in the Severn estuary in southwest England or the Bay of Fundy in Canada, the tidal amplitude will increase substantially. The Bay of Fundy, for example, has a recorded maximum tidal reach of 16.2 m, while that of the Severn estuary is 14.5 m. Note, however, that mean tidal amplitudes in these regions are likely to be much smaller than this.

The energy that can be extracted from tidal motion waxes and wanes with the tide itself. Under most conditions power output is not possible continuously. Tidal movement is, however, extremely predictable and the timing of the tides can be calculated with great accuracy. This makes tidal power a valuable form of renewable generation because, although intermittent, it is entirely reliable in its behavior.

The World Energy Council has estimated the global annual energy dissipation as a result of tidal motion to be 22,000 TWh. Of this, 200 TWh are considered economically recoverable based on the use of tidal barrages. Aggregated national estimates (see below) suggest that the total recoverable energy is much higher than this but not all of it will be economic. Current production from tidal energy is probably around 1 TWh.

There has been considerable interest in tidal power since the 1960s and a number of countries have identified sites where tidal power production would be possible. However, although a number of pilot projects have been launched, large-scale schemes have generally been judged too expensive to build.

One of the most thorough research projects examining national tidal potential was carried out in the United Kingdom between 1983 and 1994. This project looked at a range of possible schemes in England and Wales. It concluded that if every practicable tidal estuary with a spring tidal range of more than 3.5 m was exploited, around 50 TWh of power could be generated each year. This represented around 20% of the electricity consumption in England and Wales in the mid-1990s. The United Kingdom's best site is the Severn estuary. The country probably has the best tidal regime in Europe but the European Atlantic coast offers a variety of other potential sites.

In Canada, the Bay of Fundy has the highest tides in the world. This region, on Canada's east coast, has been the subject of intense examination. A comprehensive study of the region, carried out in the mid-1960s, focused on sites with a total generating capacity of nearly 5000 MW. However, tentative schemes to build projects were abandoned during the changing economic climate at the end of the 1970s.

Russia has significant potential for tidal generation, particularly in the White Sea on the Arctic coast and in the Sea of Okhotsk. A site at Mezenski Bay on the White Sea could provide 15 GW of generating capacity and an annual output of 40 TWh, while a second at Tugurki Bay has a potential generating capacity of 7800 MW and 20 TWh/year. The Russian state utility has estimated that total Russian tidal potential is 250 TWh/year.

Korea has a variety of tidal sites and is home to the world's largest tidal power plant at Sihwa that began operating in 2011. The country has plans for further large tidal schemes. India also has substantial tidal potential. The Gulf of Kutch on the northwest coast has been studied and a 600 MW project proposed. Meanwhile, the Gulf of Khambhat has an estimated generating capacity of 7000 MW. The Indian government has put the country's tidal potential at 10,000 MW but it could be much larger than this.

China has studied various potential sites. Its southeast coastline is thought to offer particularly good opportunities. Mexico has looked at a site on the Colorado estuary, Brazil and Argentina have studied projects, and the United States has examined a site in Alaska.

Australia's northwestern coast has some of the highest tidal ranges in the world and there are a number of inlets that could be harnessed to generate electricity. A novel two-basin project was proposed near the town of Derby but the scheme was rejected by the Western Australian government in 2000 in favor of a fossil fuel plant.

Tidal power does not need to be tied to estuaries. In the 1960s, France developed plans for an offshore project in Mont St. Michel Bay. The scheme was shelved when the country decided to invest heavily in nuclear power. The Mont St. Michel project involved a tidal plant that did not make use of an estuary. Instead, a circular barrage, or bund, was to be constructed that would completely enclose an area of open sea. This type of plant would operate in exactly the same way as an estuary plant, with water flowing into the enclosed reservoir when the tide rises, and flowing out through turbines during the ebb tide. While this approach would involve enormous construction costs, it does have the merit of allowing a large tidal plant to be built where no suitable estuary exists.

OPERATING TIDAL BARRAGE POWER PLANTS

Harnessing tidal motion to generate mechanical power has a long history. Tidal basins were being used in Europe to drive mills to grind grain before 1100 AD. These plants were widely replaced when the Industrial Revolution introduced steam engines and fossil fuel, but a few survived though there are none now operating commercially. The exploitation of tidal ebb and flow to generate electricity has been less well tried. Table 9.1 shows the most important tidal power plants that have been built. As the table indicates, the largest is at Sihwa in South Korea, followed closely by La Rance on the northwest coast of France close to St. Malo.

The 240 MW La Rance plant was built using specially devised bulb turbines. A small turbine of similar design was bought by the Russian government during the 1960s, but it is not known whether it was ever deployed, although there has been speculation that the 400 kW Kislaya Guba project represents the final resting place for that turbine.

TABLE 9.1 World's Tidal Power Plants

Site	Country	Capacity (MW)	Year Entered Service
Various	China	11	1958 onwards
La Rance	France	240	1966
Kislaya Guba	Russia	(0.4)	1968
Annapolis	Canada	18	1984
Sihwa	South Korea	254	2011

Source: World Energy Council, Modern Power Systems, International Waterpower and Dam Construction.

After La Rance, the third largest tidal barrage project is at Annapolis Royal on the Bay of Fundy in Canada. China has also developed some small-scale projects, of which the largest is at Jiangxia. Work on tidal power generation began in China in 1958 and there are thought to be seven projects in operation today with an aggregate capacity of 11 MW.

TIDAL POWER PLANT DESIGN

The main component of any tidal barrage power plant is the dam or barrage that is built across the mouth of a tidal estuary or inlet. This barrage is fitted with special sluice gates that are opened and closed during different stages of the tide. It is also fitted with hydropower turbines, and these too are equipped with gates so that seawater can either be allowed to pass through them or prevented from doing so as the tide changes (Figure 9.2).

The simplest and most common form of power generation with a tidal barrage is ebb-generation. Under this scheme water is allowed to pass from the sea

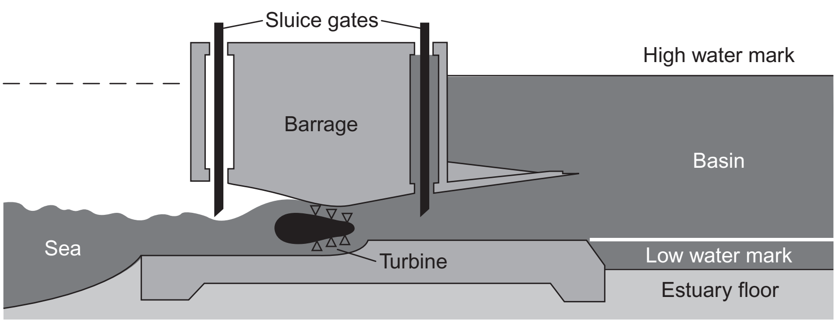

FIGURE 9.2 Cross section of a typical tidal barrage.

across the barrage and into the lagoon or basin behind it as the tide rises. Once high tide has been reached, the sluice gates within the barrage are shut, trapping the seawater within the lagoon.

The tide is now allowed to fall on the seaward side of the barrage to create a head of water across the barrage that will drive the trapped water through the hydro-turbines. The length of time during which the water is held will depend on the specific project but will normally be until the tide has fallen to around half its tidal range. At this point the gates closing the turbines are opened, allowing the water from the lagoon to flow through them and back to the sea. Generation will normally continue until close to or after low tide.

When generation stops, the gates protecting the turbines are closed again and the sluice gates opened so that as the tide turns, water will once again pass the barrage into the lagoon. The cycle is then repeated through the next tide.

It is possible to reverse the mode of operation and generate power on the flood tide instead of the ebb tide. In this case the sluice gates are kept closed at low tide so that no water can pass. When the tide has risen by about half its range the turbine gates are opened, allowing water to flow through them and into the lagoon, generating power in the process. Generation continues until levels on either side of the barrage are similar when the turbine gages are closed and the sluice gates opened, permitting the lagoon to empty again. While this is operationally simply ebb-generation in reverse, it is not commonly employed because it leaves the tidal basin behind the barrage exposed to low-tide conditions for extended periods, a situation that can have damaging environmental effects. (However, the particular conditions at the Sihwa tidal barrage plant in South Korea have made this mode of operation preferable.)

It is also possible to generate power during both the ebb and flow tide. The French plant at La Rance was designed to operate in this way but the plant in fact only operates in ebb-generation mode. The main problem with two-way generation is that calculations suggest that economic gains are small and unlikely to be cost effective because it necessitates the additional expense of either two-way turbines or two sets of turbines, one for each direction. On the other hand, it allows generation to take place for much more of the tidal cycle and leads to an overall lower peak power since the head of water that develops is never as high as with single-direction operation. This would in principle allow smaller turbines to be used.

A further operational mode, one that has been employed at La Rance in France, is to use the turbines as pumps to pump additional water across the barrage. Pumping takes place at close to high tide, creating a larger head of water than would be available from the tidal range alone. It is possible to generate up to 10% more power using pumping than without it, and the economics are attractive since the pumping takes place when the head height across the barrage is very small, therefore requiring little additional energy while energy is returned from a much higher head.

TWO-BASIN PROJECTS

A conventional one-basin tidal barrage project can only generate power during a part of each tidal cycle. To get around this a variety of two-basin projects have been proposed. This adds complexity but allows either continuous generation or generation for a longer period than a single-basin design.

One type of two-basin design comprises two single basins, each with its own barrage controlling the flow of seawater in and out (Figure 9.3). These two basins are then connected by a channel into which turbines are fitted. In operation, one of the basins opens its sluice gates only close to low tide, keeping the water level within its lagoon as low as possible. Meanwhile, the second opens its sluice gates toward high tide so that the water level within its lagoon is always high.

Water is then allowed to flow from one lagoon to the other through the channel linking the two. The flow rate and the capacity of the turbine within the channel is sized so that there is always more water in the high-water lagoon than in the low-water lagoon so that there will always be a head of water to drive the turbine.

The best developed project of this type was one proposed for construction near Derby in Western Australia. The project involved building barrages across two adjacent inlets and creating an artificial channel connecting the two basins formed by these barrages. A power station with turbines capable of generating 48 MW was to be stationed on this artificial channel. However, the project was never built.

An alternative two-basin scheme design has a primary reservoir that acts like a normal ebb-flow tidal plant, generating power on the ebb tide. However, on

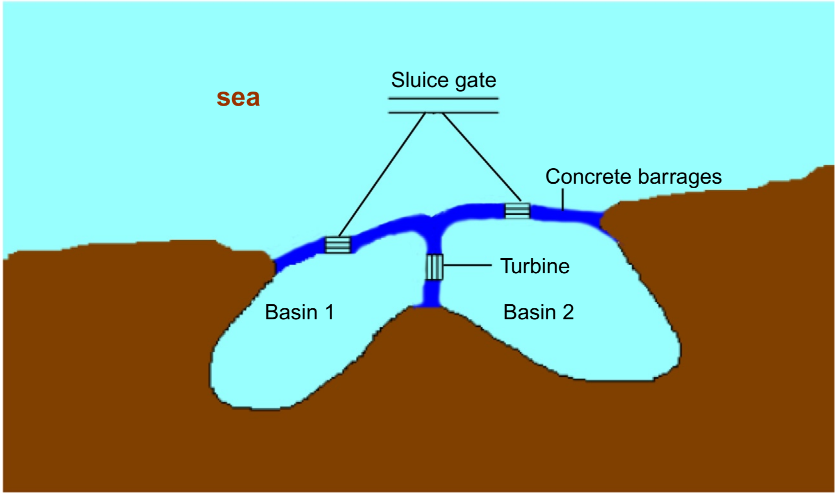

FIGURE 9.3 Layout of a two-basin tidal power project.

the seaward side of the primary basin is a second smaller basin. During the generation phase of the main basin, some power is used to pump water into the second basin, creating a storage basin from which power will always be available for generation whatever the state of the tide. The economics of such a scheme are relatively low at around 30% cycle energy efficiency.

BUNDED RESERVOIR

Instead of building a barrage across an estuary, it is theoretically possible to enclose an area of a tidal estuary or tidal region off the shoreline with an embankment or bund such as the St. Malo project discussed earlier so that it does not affect any part of the coast. The principle involved is the same, creating a reservoir that can be filled at high tide and then allowed to empty when the tide has fallen.

Although the St. Malo project never progressed beyond the design and planning stage, tidal lagoons of this type are being proposed again, together with what are proposed to be cost-effective ways of building them. Such designs have environmental advantages because they do not affect a tidal estuary or coastal land region itself. In addition, they can be built so that they do not obstruct waterways or shipping routes. Shallow tidal flats in areas of high tidal reach are judged to be the most economical sites for constructing such plants. There is some interest in such projects at the beginning of the second decade of the 21st century but no scheme of this type has so far been constructed.

TIDAL BARRAGE CONSTRUCTION TECHNIQUES

The construction of a tidal barrage represents the major cost of developing tidal power. As a result, much of the research work carried out into tidal power has focused on the most efficient way of building the barrage.

Construction of the French tidal power plant at La Rance was carried out behind temporary coffer dams, enabling the concrete structure to be built under dry conditions. While La Rance was completed successfully using this approach, the method is generally considered too expensive as a means of constructing a tidal barrage today. There is also an environmental problem attached to completely sealing an estuary for the period of construction, which might easily stretch into years. For that reason, such an approach is unlikely to be adopted for the future.

A novel approach suggested for the construction of a barrage across the River Mersey in England borrows something from the construction of La Rance. The idea proposed here was to procure a pair of redundant bulk carriers (e.g., oil tankers) and sink them on the riverbed parallel to one another, sealing the ends and filling the enclosed space with sand to create an island. Concrete construction would be carried out on the island as if it were dry land. To create a watertight structure, diaphragm walls would be fabricated of reinforced concrete; the turbines and sluice gates required for the operation of the power station would subsequently be fitted to this concrete shell.

Once the first section of the barrage had been completed the bulk carriers would be refloated, moved along to the next section, and sunk again. This process would be repeated, until the barrage had been completed. The River Mersey barrage has not been built, so the efficacy of the method has yet to be tested.

Where an estuary is shallow, an embankment dam could be constructed instead of a concrete dam using sand and rock as its main components. Sand alone would not make a stable embankment; wave erosion would soon destroy it. Therefore, some form of rock reinforcement would be required on the seaward side. Concrete faces on both sides of the embankment could provide further protection. The sand needed for construction of such an embankment might be recovered from the estuary by dredging. Rock could also be removed from the riverbed by blasting or brought to the site from elsewhere. Rock is a more expensive construction material than sand so its use would have to be minimized to keep costs as low as possible.

While all these methods have their attractions, the construction method most likely to be used to build a large barrage today would involve prefabricated units called caissons. Made from steel or concrete, the caissons would be built in a shipyard and then towed to the barrage site where they would be sunk and fixed into position with rock anchors and ballast. Some caissons would be designed to hold turbines, others would be designed as sluice gates, and a third type would be blank. These would be placed between the other two types to complete the barrage.

Caisson construction was the favored approach in a study for construction of the Severn barrage in England completed in 1989 under the auspices of the Severn Barrage Development Project. A turbine caisson for this project would have weighed over 90,000 tons and would have had a draft of 22 m. The minimum height of the vertical faces would be 60 m. As a result of their size, special facilities would have been needed to construct them. Prefabrication of the caissons was expected to reduce construction time to a minimum. Even so, the Severn project was scheduled to take 10 years to complete.

For offshore lagoon construction, one company that is promoting such schemes proposes using a conventional rubble-mound breakwater. Barrage failure would have minimal safety consequences in an offshore lagoon because the project is self-contained, so this relatively cheap means of construction should be possible.

TURBINES

The turbines in a tidal power station must operate under a variable, low head of water. The highest global tidal reach, in the Bay of Fundy in Canada, is 15.8 m and the mean tidal reach probably half of this range; most plants would have to operate with much lower heads than this. Such low heads necessitate the use of a propeller turbine, the turbine type best suited for low-head operation. The fact that the head varies appreciably during the tidal cycle means that a fixed-blade turbine will not be operating under its most efficient conditions during the majority of the tidal flow; consequently, a variable-blade Kaplan turbine is usually employed. As is the case with most low-head hydropower plants, tidal power plants usually employ a series of small turbines running along the barrage since these can exploit the available energy more effectively than a small number of large turbines.

The most compact and efficient design of propeller turbine for low-head applications is the bulb turbine in which the generator attached to the turbine shaft is housed in a watertight pod, or bulb, directly behind the turbine runner (Figure 9.4). The whole turbine–generator assembly is then hung inside a chamber that channels the water flow through the turbine blades to extract the maximum energy possible. The La Rance tidal plant employs 24 bulb turbines, each fitted with a Kaplan runner and a 10 MW generator. Bulb turbines were new when La Rance was built and construction of the plant involved some experimental work; of the 24 turbines, 12 had steel runners and 12 had aluminum bronze runners. Experience has led the operators to prefer the steel variety.

The turbines at La Rance were designed to pump water from the sea into the reservoir behind the barrage at high tide to increase efficiency. This was found to cause severe strain on parts of the generator and the design had to be modified. Work was carried out between 1975 and 1982. Since then the plant has operated smoothly and with high availability.

The more recent Sihwa tidal plant in South Korea also uses bulb turbines. In this case the plant is equipped with ten 26 MW bulb turbines with variable-blade propeller units. This power plant is built into a sea wall erected in 1994 to create an inland lagoon where water was collected for irrigation. Since

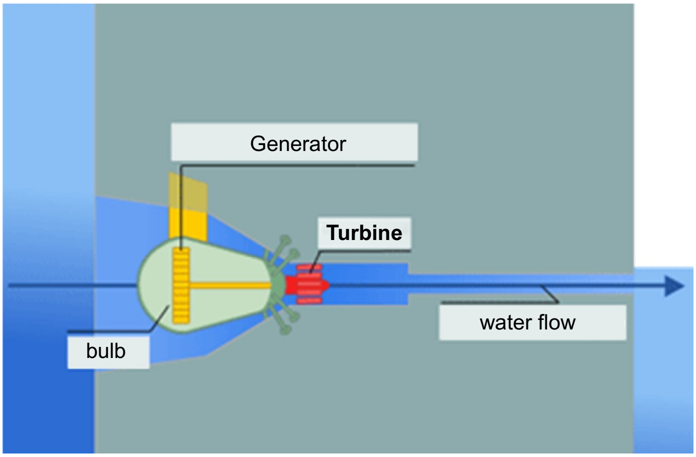

FIGURE 9.4 Cross-section of a bulb turbine in a tidal barrage.

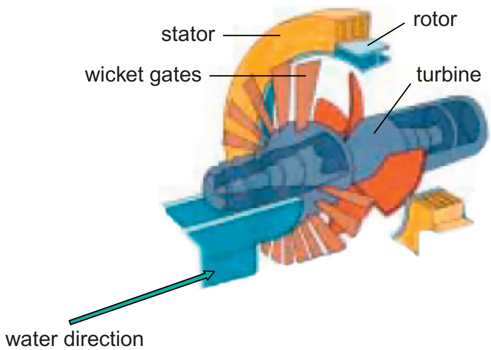

FIGURE 9.5 Cut-away of a Straflo turbine.

then industrial pollution of the lagoon has made the water unusable. The tidal plant forms part of a scheme intended to flush the lagoon to reduce pollution levels. Unlike most other tidal barrage plants, it generates on the flow tide.

An alternative to the bulb turbine is a design called the Straflo turbine (Figure 9.5). This is unique in that the generator is built into the rim of the turbine runner, allowing the unit to operate in low-head conditions while keeping most of the generator components out of the water. A single large Straflo turbine generator was installed at the Annapolis tidal power plant at Annapolis Royal in the Bay of Fundy, Canada. This 18 MW unit is the only one of similar size that has been built, so experience with the design in limited.

TURBINE SPEED REGULATION

The speed of a conventional turbine generator has to be closely regulated so that it is synchronized with the electrical transmission system to which it is attached. To aid frequency regulation under the variable conditions of a tidal power plant, a set of fixed blades called a regulator are often placed in front of the turbine blades to impart a rotary motion to the water. The use of these blades in conjunction with a variable-blade Kaplan turbine provides a considerable measure of control over the runner speed.

In small applications where such tight speed control may not be essential and where costs are critical it may be possible to use one method of control—either a variable-blade turbine or a regulator—rather than both. An isolated unit that does not connect into the grid could operate without regulation.

An alternative option is to use a variable-speed generator. This electronic solution will permit the turbine to run at its optimum speed under all conditions while delivering power at the correct frequency to the grid. This allows some efficiency gains. However, the solution is more costly than a conventional generator with mechanical speed control of the turbine. Variable-speed generators are being used on some hydropower schemes today (see Chapter 8). Capacity is limited but that is unlikely to be a problem with a tidal power plant where unit size is generally small.

SLUICES AND SHIPLOCKS

The sluices in a tidal barrage must be large and efficient enough to allow the tidal basin behind the barrage to fill with water quickly. Unless the water level behind the barrage effectively follows that on the seaward side, the efficiency of the plant is reduced. Where the water is sufficiently deep, efficient sluices can be built using the concept of the Venturi tube. Such a design will transfer water through the barrage extremely efficiently but it must be completely submerged. More conventional sluice gates usually need to be larger than Venturi tubes to provide the same rate of transfer.

Many of the rivers suitable for tidal development carry significant ship-borne trade and water traffic. To enable ships and boats to continue to use a river, ship locks must be included in the barrage. There must also be facilities to allow fish and other forms of marine life to pass the barrage. This is particularly important if the river is one used by migratory fish such as salmon.

ENVIRONMENTAL CONSIDERATIONS

Construction of a barrage across a tidal river is bound to affect the conditions on both sides of the structure. Water movement patterns will be changed, sedimentation movement will be affected, and the conditions at the margins of the estuary on both the landward and seaward side of the barrage will be altered. For a barrage across an estuary, the movement of marine animals is likely to be restricted too. This could have a dramatic effect on both marine and avian life.

The major effect of the barrage will be on water levels and water movement. Water levels will be altered on both sides of the barrage and the tidal reach may change behind the barrage, although the effect will be reduced as the distance from the barrage increases. Some areas that were regularly exposed at low tide will be continuously under water after the barrage is constructed. Though the volume of water flowing down the river should remain the same, patterns of movement will be changed.

Sedimentation will be affected in complex ways. The tidal waters of an estuary frequently bear a great deal of sediment. Some is brought in from the sea, some carried downstream by the river. Changes in current speeds and patterns caused by the interpolation of a barrage will affect the amount of sediment carried by the water and the pattern of its deposition. This will, in turn, affect the ecosystems that depend on the sediment.

Other areas of concern involve animal species. The effect on fish, particularly migratory species, is significant. Fish gates can be built to permit species to cross the barrage. Many can also pass through the sluice gates. However, there is a danger that fish will pass through the turbines too, being injured in the process. Various methods have been explored to discourage fish from the vicinity of the turbines, with patchy success.

Many species of birds live on mud flats in estuaries. There is a possibility that such mud flats would disappear after a barrage had been built, and with them the birds whose habitat they formed. Salt marshes adjacent to estuaries are also likely to be affected. Studies have been conducted at potential U.K. barrage sites to try and estimate the scale of such effects but much work remains to be done in this area.

The effects of a tidal lagoon or bunded reservoir are likely to be less dramatic than those associated with a coastal barrage. Since none has yet been built, the range of effects is currently unknown.

Elsewhere, global experience with tidal power plants is limited. What evidence there is suggests that such projects have no major detrimental effect on the environment. The evidence from La Rance, in particular, has provided no serious cause for alarm. Even so, it would be dangerous to make any assumptions. An extremely careful environmental impact assessment would form a vital part of any future tidal project.

It is important to remember when considering tidal barrage and similar projects that while these projects change the local environment, they do not destroy it. As with hydropower projects discussed in Chapter 8 there will be ways of mitigating the effects, and the new environment created after barrage construction may be as rich as the one it replaces. However, any tidal project is likely to provoke fierce debate.

COST OF ELECTRICITY GENERATION FROM TIDAL BARRAGE POWER PLANTS

A tidal barrage power plant will generally be a highly capital-intensive project and the cost of electricity during the period associated with repayment of any loan associated with the project will be higher than for a similarly sized hydropower plant, because the tidal power plant cannot operate continuously but only during certain periods of the tidal cycle.

While the high initial capital outlay is likely to provide a disincentive for many project developers, the longevity of a tidal barrage power plant means that once the loan has been repaid the plant will enjoy a long life providing low-cost electricity. The tidal power plant at La Rance, France, has been operating for over 40 years and is likely to continue for another 40 years while the electricity produced by the plant is cheaper than French nuclear power.

The most recently constructed tidal power plant is the Sihwa plant in South Korea. This 260 MW project is estimated to have cost $250 m, or just under $1000/kW. However, the power plant was installed into an existing barrage so the costs for this scheme cannot be considered typical.

One of the most intensely studied barrage projects is that on the Severn River in the United Kingdom. The latest study was carried out by the U.K. government and published in 2010.[1] This study examined a number of possible schemes across the river, concluding that the largest, with an installed capacity of 8640 MW, would have a cost between £23.2 bn and £34.3 bn or £2690–3970/kW. This figure was disputed by a consortium that claimed the barrage could be built for £17–£18 bn, a unit cost of around £2000/kW. The lifetime of the plant was estimated in the U.K. government report to be 120 years. While the initial cost of a barrage will be large, it can be partially mitigated if the structure can also provide a road or rail crossing. In the case of the Severn barrage, where no transport crossing was proposed, the U.K. government suggested that the cost of electricity from the plant would be more expensive than offshore wind or nuclear power.

Another proposed large tidal barrage project is at Incheon, South Korea. This 1320 MW scheme has an estimated cost of \$3.4 bn, or \$2580/kW. Work on the project is currently scheduled to start in 2014 with completion in 2019.

Tidal lagoons have been claimed to be cheaper to build than barrages but opinions differ. An independent review of a proposed 60 MW tidal lagoon project in the same region as the proposed Severn barrage in the United Kingdom put the cost of this project in 2006 at £234 m, or £3900/kW. This was significantly higher than the cost proposed by the company that had developed the scheme of £82 m or £1370/kW.

The preceding costs are higher than typical costs for hydropower plants of a similar size. Even so they could be economically viable. In most cases, however, such projects will only be built if they can serve additional uses and so attract public sector investment to help support the financing.

1. Severn Tidal Power: Feasibility Study Conclusions and Summary Report, U.K. Department of Energy and Climate Change, 2010.

Chapter 10

Power System Energy Storage Technologies

Energy storage plays a vital part in the modern global economy. At a national level, oil and gas are regularly stored by both utilities and governments, while at a smaller scale petrol stations store gasoline and all cars carry a storage tank to provide them with the ability to travel a significant distance between refueling stops. Domestic storage of hot water is also usual in modern homes. Yet when it comes to electrical energy, storage on anything but a small scale in batteries is rare.

Part of the reason for this is that storage of electricity, although it can be achieved in a number of ways, is difficult. In most storage technologies the electricity must be converted into some other form of energy before it can be stored. For example, in a battery it is converted into chemical energy, while in a pumped storage hydropower plant the electrical energy is turned into the potential energy contained within an elevated mass of water. Energy conversion makes the storage process complex and the conversion itself is often inefficient. These and other factors help to make an energy storage system costly.

In spite of such obstacles, large-scale energy storage plants have been built in many countries. In the majority of cases these installations are pumped storage hydropower plants, often built to capture and store power from base-load nuclear power plants during off-peak periods. Many of these storage plants were built in the 1970s. More recently there has been renewed interest in technologies such as pumped storage for grid support, particularly in European countries that are installing large capacities of renewable capacity such as wind and solar power. However, the economics of energy storage often make construction difficult to justify in a liberalized electricity market.

While economics may not always favor their construction, energy storage plants offer significant benefits for the generation, distribution, and use of electric power. At the utility level, for example, a large energy storage facility can be used to store electricity generated during off-peak periods (typically overnight), and this energy can be delivered during peak periods of demand when the marginal cost of generating additional power can be several times the off-peak cost. Energy arbitrage of this type is potentially a lucrative source of revenue for storage plant operators and is how most pumped storage plants operate.

At a smaller scale, energy storage plants can supply emergency backup in case of power plant failure, as well as other grid support features that help to maintain grid stability. They can also be employed in factories or offices to take over in case of a power failure. Indeed, in a critical facility where an instantaneous response to loss of power is needed, a storage technology may be the only way to ensure complete reliability.

Energy storage also has an important role to play in the efficient use of electricity from renewable energy. Many renewable sources of energy, such as solar, wind, and tidal energy, are intermittent and so are incapable of supplying electrical power continuously. Combining some form of energy storage with a renewable energy source helps remove this uncertainty and increases the value of the electricity generated. It also allows all the renewable energy available to be used. Today, the shedding of excess renewable power when demand does not exist for it, or when the grid cannot cope with it, is becoming common on some grid systems with high renewable capacity.

While there are many types of electrical energy storage systems, pumped storage hydropower plants account for virtually all grid storage capacity available today with perhaps 130 GW of generating capacity in operation, based on estimates by the International Hydropower Association.[1] This was effectively the only large-scale energy storage technology available until the late 1970s but in the past 30–40 years new interest has been stimulated and a range of other technologies have been developed. These vary in size so that some are suitable for transmission system–level storage while others are more suited to the distribution grid or even for small micro-grids. They include a range of battery storage systems, compressed-air energy storage, large storage capacitors and flywheels, superconducting magnetic energy storage, and systems designed to generate hydrogen as an energy storage medium. The widespread adoption of electric vehicles that use battery energy storage could potentially offer a major new means of storing grid electricity too.

If deployed widely, these technologies could potentially transform the way the grid-based delivery of electrical energy operates by eliminating the need for expensive peak power plants while at the same time integrating the range of renewable generation technologies now available. This would, in turn, eliminate the vulnerability of electricity production to the vagaries of the global fossil fuel markets, creating more stable economic conditions everywhere. There is no consensus on how much storage capacity would be required to achieve this on a mature national grid but it could be equivalent to around 10–15% of the available generating capacity.

In spite of the apparent advantages offered by energy storage, widespread adoption remains slow. Cost appears to be the main obstacle, although developments are slowly bringing costs down. At the same time the growth of

1. *2013 IHA Hydropower Report*, International Hydropower Association.

distributed generation is offering new opportunities for small-scale energy storage facilities. This chapter will look at the range of technologies available and where they might fit into the electricity system.

TYPES OF ENERGY STORAGE

Electricity is an ephemeral form of energy that normally has to be used as soon as it has been generated. This is why the role of system operators and their electricity dispatching systems are important; they have to balance the demand for electricity with its supply. If one fails to match the other, problems occur: system voltages rise or fall and grid frequency may change and cause problems across the grid. It would seem obvious, given this situation, that some reservoir of saved electricity would be a major boon to grid operation. Yet storing electricity has proved difficult to master.

Storing electricity in its dynamic form, amps and volts, is almost impossible. The nearest one can get is a superconducting magnetic energy storage ring that will store a circulating DC current indefinitely provided it is kept cold. A capacitor storage system stores electricity in the form of static electric charge. All other types of energy storage convert the electricity into another form of energy. This means that the energy must then be converted back into electricity when it is needed.

A rechargeable battery may appear to store electricity but in fact it stores the energy in chemical form. A pumped storage hydropower plant stores potential energy, a flywheel stores kinetic energy, and a compressed-air energy storage plant stores energy in the form of compressed air, another type of potential energy. Alternatively, one might use electrolysis to turn electricity into hydrogen, yet another chemical transformation of the energy. This can then be burned in a thermal power plant or used in a fuel cell to turn it back into electricity.

The storage systems based on these processes are all viable ways of storing electricity and many of them are commercially available. Each has different characteristics, such as response time and storage efficiency, that helps differentiate the technologies and define their applications.

Some of these systems can deliver power extremely rapidly. A capacitor can provide power in 5 ms, as can a superconducting energy storage system. Flywheels are very fast too, and batteries should respond in tens of milliseconds. A compressed-air energy storage plant probably takes 2–3 minutes to provide full power. Response times of pumped storage hydropower plants can vary between 10 seconds and 15 minutes. This technology is generally suitable for peak power delivery but less suited to fast-response grid support.

The length of time the energy must be stored will also affect the technology choice. For very long-term storage of days or weeks, a mechanical storage system is best, and pumped storage hydropower is the most effective provided water loss is managed carefully. Batteries are also capable of holding their charge for extended periods. Energy loss in other systems will make them less practical for long-term storage. For daily cycling of energy, both pumped storage and compressed-air energy storage are suitable, while batteries can be used to store energy for periods of hours. Capacitors, flywheels, and superconducting magnetic energy storage are generally suited to short-term energy storage, though flywheels can be used for more extended energy storage too.

Another important consideration is the efficiency of the energy conversion process. An energy storage system utilizes two complementary processes: storing the electricity and then retrieving it. Each will involve some loss. The roundtrip efficiency is the percentage of the electricity sent for storage that actually reappears as electricity again. Typical practical figures for different types of systems are shown in Table 10.1.

Electronic storage systems such as capacitors can be very efficient, as can batteries. However, the efficiencies of both will fall with time due to energy leakage. Flow batteries, where the chemical reactants are separated, perform better in this respect and will maintain their roundtrip efficiency better over time. Mechanical storage systems, such as flywheels, compressed-air energy storage, and pumped storage hydropower, are relatively less efficient. However, the latter two, in particular, can store their energy for long periods if necessary without significant loss.

All these factors must be taken into consideration when considering the most suitable energy storage technology for a given application. For large-scale utility energy storage there are three possible technologies to choose among: pumped storage hydropower, compressed-air energy storage, and, at the low end of the capacity range, large batteries. Batteries can also be used for small- to medium-size distributed energy storage facilities,[2] along with flywheels and

TABLE 10.1 Roundtrip Efficiency of Energy Storage Technologies

Energy Storage Technology	Roundtrip Storage Efficiency
Capacitors	90%
Superconducting magnetic energy storage	90%
Flow batteries	90%
Compressed-air energy storage	65%
Flywheels	80%
Pumped storage hydropower	75–80%
Batteries	75–90%

2. Distributed storage facilities may be used by utilities to improve local grid stability or they may be used by consumers to make their own supplies more secure.

capacitor storage systems. Meanwhile, fast-acting, small superconducting magnetic energy storage units are being used to aid grid stability. Superconducting facilities have been considered in the past for large-scale energy storage too, but they appear to be prodigiously expensive based on the technology available today.

PUMPED STORAGE HYDROPOWER

Pumped storage hydropower is both the simplest and most widely used technique for storing electrical energy today. It was first deployed in Switzerland around 1904,[3] and there is probably around 130 GW of capacity in use, though estimates vary. These plants vary in size from a few megawatts to over 1000 MW, with the largest close to 3000 MW in capacity. Plants can be found in Australia and China across Europe and in Russia, but the largest aggregate capacities are in Japan and the United States. Many are used in conjunction with nuclear power plants so that the latter can operate at full power irrespective of demand. However, some smaller plants are also used for peak shaving and load management duties independent of the availability of nuclear power.

The reservoir-based pumped storage plant is an adaptation of the conventional hydropower plant to enable it to operate reversibly (Figure 10.1). In a conventional hydropower plant with a reservoir water collects in the reservoir and is then released through the plant's turbines as power is required. This confers an element of energy storage but the water in the reservoir can only be used once.

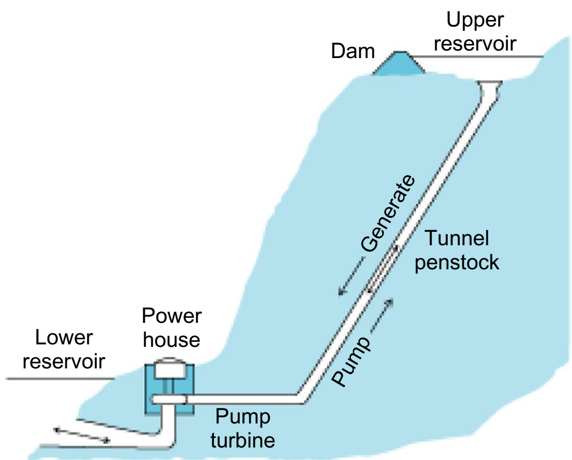

FIGURE 10.1 Pumped storage hydropower plant.

3. The earliest recorded pumped storage power plant was a 1.5 MW facility at Schaffhausen in Switzerland.

In a pumped storage plant there is a second reservoir below the turbine hall. Water that has been released from the first reservoir and used to generate electricity is collected again here and stored. The power station is also equipped with pumps—in most cases its power turbines can be operated in reverse as pumps—and during periods when excess power is available on the grid these pumps are used to pump the water that has been collected in the reservoir below the turbine hall back into the higher reservoir above the turbine hall. The water is cycled between the two reservoirs to provide either power or energy storage as needed.

This type of plant is extremely robust and though roundtrip efficiency is lower than for some other technologies, long-term energy losses are low. Leakage and evaporation are the main sources of loss and, if these are managed well, water loss can be kept small. Today this is the only technology available for very large-scale energy storage.

Pumped Storage Technology

The basic layout of a pumped storage hydropower plant involves two reservoirs, one above the other, and a turbine/pumping hall capable of both generating power from the stored water in the upper reservoir and pumping water from the lower reservoir back to the upper. For hydropower plants, in general, the energy available from a given volume of water is greater, the greater the head of water. In the case of the pumped storage plant this head is the vertical distance between the upper reservoir and the turbines. The greater this distance, the more energy a given quantity of water can store; put another way, the larger the head, the smaller the volume of water needed for a given amount of energy. However, the pumped storage head will be limited by the type of turbine that can be utilized.

While the highest head available would in theory be best, very high heads require Pelton turbines to exploit them efficiently and these cannot be used as pumps. A very high-head plant would, therefore, require separate pumps and turbines, as was used in the earliest pumped storage facilities. Using separate pumps and turbines is more expensive than using a single-pump/turbine unit. Therefore, most pumped storage plants use Francis or Deriaz turbines, which can be used in both modes. This limits the head that can be used to achieve good efficiency to about 700 m. Modern multistage pump turbines may be capable of extending this to around 1200 m.

A pump turbine may not achieve the efficiency possible when using independently optimized pumps and turbines, but the best combined pump turbines are capable of reaching around 95% efficiency for generation and 90% for pumping, leading to a roundtrip efficiency of 86%. Most plants operate in the 75–80% range.

The Francis turbine is the type most commonly used as a pump turbine. While highly efficient, it has fixed blades so the blade design is generally a compromise designed to optimize both generation and pumping efficiency. The Deriaz turbine is of similar design but with adjustable blades making it possible to optimize for generation and pumping independently. These have been used for pumped storage plants in several parts of the world but will generally be more costly than Francis turbines because of the additional complexity.

Variable-Speed Operation

While most pumped storage hydropower plants have been built using pump turbines that operate at a fixed speed that is synchronized with the grid there are significant advantages to be gained by having the ability to operate at variable speed. For pumping, variable-speed operation allows the pump to function with surplus power at different demand levels and it can take power while the demand level is changing on the grid, allowing for much greater flexibility of operation. In generation mode, variable-speed operation allows the unit to supply varying quantities of power. A fixed-speed turbine can only supply its rated output at grid frequency.

Variable-speed operation requires that the turbine generator be decoupled from the grid through a power electronic interface so that the electricity produced by the turbine at variable frequency can be injected into the grid at the synchronous frequency, or the latter can be adjusted when in pumping mode to allow variable pumping capacity. Therefore, a variable-speed pump generator will be more costly than a fixed-speed unit.

Pumped Storage Sites

The single greatest limitation (aside from economics) faced by pumped storage hydropower is the availability of suitable sites. A plant of this type requires two reservoirs at different heights. This can be difficult to engineer.

In rare cases it is possible to find two existing lakes that can be utilized to create a pumped storage facility. If natural lakes are exploited, it will be necessary to take into account the fact that the water level in both will vary more widely than it would naturally and assessing the environmental impact of these changes. More commonly, a natural lake might form one reservoir while the second is human-made. The third option is for both reservoirs to be human-made. However, this can add significantly to the capital cost of such a plant.

A further, and yet rarely attempted, solution is to use the sea as the second or lower reservoir. This is the layout of the 30 MW Yanbaru seawater pumped storage plant in Japan. One other idea that has never yet been exploited is to bury the lower reservoir underground in a suitable geological formation.

Plant capability depends on both the size of the reservoirs and the head, or vertical, height between them. The volume of the reservoirs will determine the overall capacity of the plant to store and supply energy. The more water, the more energy it can contain. However, for a given storage capacity, the output will depend both on the size of the turbines and the head. A high head can deliver more power from a given flow of water than a small head.

Performance

Pumped storage plants are capable of rapid response to sudden changes in demand. The Dinorwig plant in Wales can go from standby to synchronization at full output of 1320 MW in 12 seconds, a ramp rate of 110 MW/sec. Other, more modern pumped storage plants have been specified to be capable of output ramping when in operation at up to 500 MW/min, which compares favorably with gas turbine peaking plants and modern combined cycle power plants.

With adequate control of evaporation and leakage, the energy stored in a pumped storage reservoir can be retained indefinitely. This is not a significant advantage when most plants cycle daily, but it could prove so under circumstances where energy needs to be stored seasonally, such as solar power for use in winter, or wind power for use in less windy summer months.

Costs

An energy storage plant such as a pumped storage hydropower plant will depend for its revenue on being able to buy power at low cost and then sell it at higher cost. The income will, therefore, vary depending on a wide range of conditions. From an economic point of view the capital cost of building the plant will be the most important factor in determining its viability. This is likely to be relatively high because, like most hydropower plants, pumped storage is a capital-intensive technology.

At the top end, capital costs are likely to be as high or higher than for a traditional hydropower plant, which, as shown in Chapter 8, are generally in the range \$1000–2000/kW. A 500 MW plant proposed for construction in California has an estimated cost of \$1.1 billion and a capacity of 500 MW, or around \$2200/kW. In contrast, the Tianghuangping pumped storage plant in Zhejiang province, China, cost \$1.1 billion for 1800 MW when it came on-line in 2001, around \$600/kW. Much of the difference can probably be accounted for by the lower labor costs in China.

Small pumped storage plants are likely to be relatively more expensive than larger installations.

COMPRESSED-AIR ENERGY STORAGE

Compressed-air energy storage (CAES) is a system whereby energy is stored in the form of air pressurized above atmospheric pressure. Compressed air has a long history as a means of both storing and distributing energy. Systems based on this energy distribution medium were installed during the late 19th century in cities as various as Paris (France), Birmingham (U.K.), Dresden (Germany), and

TABLE 10.2 Commercial CAES Plants

Plant	Generating Capacity	Commissioning Date
Huntorf, Germany	290	1987
McIntosh, Alabama, U.S.	110	1991

Buenos Aries (Argentina) to supply power for industrial motors and commercial use in a variety of applications including the textile and printing industries.

The use of compressed-air storage as an adjunct to the power grid began with the construction of the Huntorf power plant that was built in Germany in 1978 but only operated commercially for 10 years. A second CAES plant was built by the Alabama Electric Cooperative in the United States and entered service in 1991. The latter facility has continued to provide storage services ever since. Details of these two plants are shown in Table 10.2.

In spite of being championed by organizations such as the U.S. Electric Power Research Institute, no further commercial project has ever been built, although others have been proposed and work even started on two. Even so, CAES remains of interest because it is the only other very large-scale energy storage system after pumped storage hydropower. Individual CAES plants are generally smaller than typical pumped storage plants but sites suitable for their construction are much more widespread than those for the hydropower storage plants. They could, therefore, provide a more widely distributed large-scale energy storage network.

Compressed-Air Energy Storage Principle

A CAES plant requires two principal components: a storage vessel in which compressed air can be stored without loss of pressure, and a compressor/expander to charge the storage vessel and then extract the energy again (Figure 10.1). (The latter might, in fact, be a compressor and a separate expander.) In operation the plant is broadly analogous to the pumped storage hydropower plant. Surplus electricity is used to compress air with the compressor and the higher-pressure air is stored within the storage chamber. This stored energy can then be retrieved by allowing it to escape through the expander, an air turbine that is essentially a compressor operating in reverse. The expanding air drives the air turbine, which turns a generator to provide electrical power.

The compression and expansion functions of the CAES plant can be performed by the two primary components of a standard gas turbine. As seen in Chapter 4, the gas turbine comprises three components: a compressor, combustion chamber, and turbine. If the combustion chamber is removed and the two rotary components separated, then these can alternately use electricity to

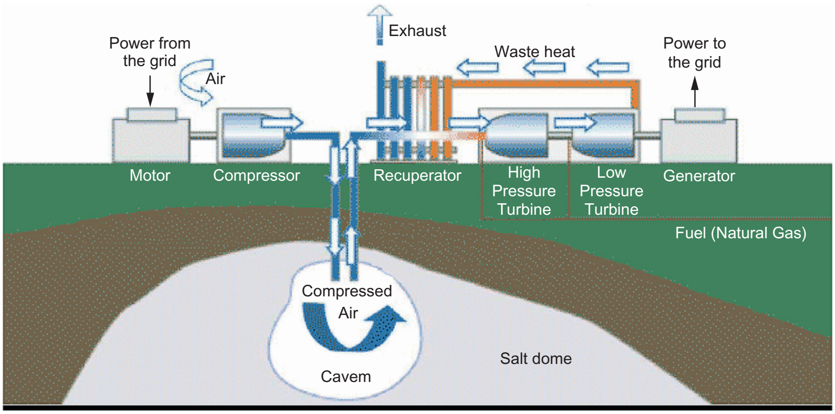

FIGURE 10.2 A compressed-air energy storage plant.

compress air for storage and extract energy from it again to regenerate the stored power (Figure 10.2).

In practice, a slightly different arrangement is preferred that is closer still to the gas turbine. A rotary gas turbine compressor is used to compress air that is then stored in the storage chamber. When power is required, compressed air is extracted again and fed into a combustion chamber where it is mixed with fuel and ignited, generating a higher-pressure, higher-temperature thermodynamic fluid that is then used to drive the turbine stage of the plant.

Since a plant operated in this way requires natural gas or another fuel, it is not a straightforward energy storage system. However, the economics of this mode of operation appear to be the most attractive because it can generate more electricity than was used to store the compressed air. Additional generation is between 25% and 60% depending on the plant design. A further advantage is that the turbine stage of the plant does not have to drive the compressor as it would in a conventional gas turbine, so it can generate up to three times more power than it would when coupled to a compressor. Therefore, turbines for CAES plants are relatively smaller than for a similar generating capacity gas turbine.

Compressed-Air Storage Facilities

The most important part of a CAES plant is somewhere to store the compressed air. Small-scale CAES plants—with storage capacities up to 100 MWh and outputs up to 20 MW—can use aboveground storage tanks built with steel pressure vessels but large, utility-scale plants need underground caverns in which to store the air. The natural gas industry has used underground storage caverns for years to store gas; these same caverns can provide ideal storage facilities for a CAES plant. However, the demand for such a cavern limits the development of CAES to places where such storage caverns are available.

A number of different types of underground caverns can be exploited. The most expensive is a human-made rock cavern excavated in hard rock or created by expanding existing underground mine workings. Such a site must be located in an impervious rock formation if it is to retain the compressed air without loss so the suitability of underground coal mines and limestone mines will depend on whether they are air-tight.

Salt caverns are another type of storage site, one that has been commonly used for gas storage. These are created within naturally occurring underground salt domes by drilling into the dome and pumping in water to dissolve and remove the salt to create an enclosure. Salt deposits suitable for such caverns occur in many parts of the world.

It is another type of geological structure, however, an underground porous rock formation, which offers the cheapest underground storage facility. Structures of this type suitable for gas storage are found where a layer of porous rock is covered by an impervious rock barrier. Examples can be found in water-bearing aquifers, or in porous underground strata from which oil or gas have been extracted. Aquifers can be particularly attractive as storage media because the compressed air will displace water within the porous rock, setting up a constant-pressure storage system. With rock and salt caverns, in contrast, the pressure of the air will vary as more is added or released.

All three types of underground storage structures require sound rock formations to prevent the air from escaping. They also need to be sufficiently deep and strong to withstand the pressures imposed on them. It is important, particularly in porous-rock storage systems, that there are no minerals present that can deplete the oxygen in the air by reacting with it. Otherwise, the ability of the air to react with the fuel during combustion will be affected, reducing the power available during the generation phase of the storage generation cycle.

Underground rock structures capable of storing compressed air are often widely available. For example, a survey in the United States found accessible sites of different types across 80% of the country.

Turbine Technology and CAES Cycles

A CAES plant generally exploits standard gas turbine compressor and turbine technology, but because the two units operate independently, they can be sized differently to match the requirements of the plant. The larger the compressor compared to the turbine, the less time it requires to charge the cavern with a given amount of energy. The Hundorf plant that was built in Germany required 4 hours of compression to provide 1 hour of power generation, whereas the McIntosh plant in Alabama needs only 1.7 hours of compression for 1 hour of generation.

As a consequence of compression and generation being separated, a CAES plant turbine can operate well at part load as well as full load. More complex operation is also possible. The Alabama plant, for example, uses two turbine stages with the exhaust from the last turbine is used to heat air from the cavern before it enters the first turbine. Fuel is not actually burned in the compressed air until it enters a combustion chamber between the first and second chambers.

Many of the refinements used to improve gas turbine performance outlined in Chapter 4 can be used in CAES plants too. For example, the compressor can be divided into two sections with air cooling between the stages to reduce its volume (intercooling) and heat from the turbine exhaust can be recovered and used to heat the compressed air extracted from the storage chamber (recuperation). Reheating, where the turbine is divided into two stages with an additional combustion chamber between stage one and stage two can also be applied versions of these latter two are used in the McIntosh plant, as noted before.

Mechanical components are never 100% efficient so there are consequential energy losses during compression and expansion in a CAES plant. There is also an additional source of energy loss. When air is compressed it generates heat and this heat energy is lost in a conventional CAES plant. A proposed additional refinement to the conventional mode of operation involves capturing this heat and storing it for use to heat the pressured air as it exits the storage chamber, before entering the turbine. This adiabatic cycle could theoretically be used to design a CAES plant that has no need for additional fossil fuel and that could achieve a roundtrip efficiency of 65%.

In principle a CAES plant could be of virtually any size and one proposed project would have had a generating capacity of 2400 MW. However in practice most schemes are likely to be smaller than this, in the tens or hundreds of megawatts range. Startup for the two plants that have operated was around 12 min but both could be brought into service in 5 min if necessary. Round trip efficiency without the use of additional fuel will be low for conventional CAES plants such as the two that have operated but, as noted, refinements could improve this.

Costs

There is little experience with CAES so any cost estimates must be considered tentative. However, it would appear to be an economically attractive option for energy storage. Proposals within the last 10 years or so for conventional CAES plants in the United States have had installed costs of $400–900/kW depending on size and storage type. An adiabatic plant is likely to be much more expensive, with potential costs as high as $1700/kW.

LARGE-SCALE BATTERIES

Batteries are the most widely known means of storing electrical energy. Invented during the 19th century, they are now used for a whole range of portable applications from starter motor power in vehicles to providing an electrical source for mobile phones, tablet computers, and tiny electronic devices such as hearing aids. More recently, batteries have also been used in a range of mostly small renewable energy applications, and in addition some large cells have been used for grid storage and stability uses.

All batteries are electrochemical devices that convert the energy that is released during a variety of chemical reactions into electrical energy. This energy would normally be released as heat if the reaction was permitted to proceed conventionally by mixing the reactants. In an electrochemical cell (i.e., a battery) the reaction is controlled in such a way that most of this heat can be converted into electricity.

There are two distinct battery types in common use: primary cell and secondary cell. A primary cell (or battery) can only be used once. After that it cannot be recharged. However, a secondary cell is capable of being recharged, reversing the internal chemical reaction and regenerating the reactants that provided the power in the first place. It is these secondary cells that are of use in the power and utility industries.

Secondary cells can be further divided into two types: standard secondary cells and flow batteries. Standard cells are the type found in portable computers or vehicles. They are completely self-contained, and have no mechanical parts. Charging and discharging is carried out via the cell terminals and all of the reactants required are contained within the battery package. Secondary cells can be found in two types: shallow discharge cells, such as those used for vehicle starter power, which are never fully discharged, and deep discharge cells that can be completely exhausted without damage.

A flow battery is different because the actual cell within which the chemical reactants react and generate electricity does not carry the reactants themselves. Instead, these are stored in external reservoirs and pumped through the cell as required. This type of battery is more complex than a conventional secondary cell, but it has the advantage that battery capacity is limited by reservoir size and this can easily be increased for relatively little cost. On the other hand, flow batteries tend only to be economical in large sizes because of the additional cost associated with their construction.

Battery Principle

A battery is a device that can exploit a chemical reaction to produce electricity. The reactions upon which a cell is based will define the particular cell type. In all cases the reaction will occur spontaneously if the reactants are mixed, generating heat in the process. However, in the battery the reactants are separated and only allowed to react in a particular way.

The chemical reaction used in every battery can notionally be divided into two half-reactions and the battery will contain two electrodes, called the anode and cathode; each electrode is associated with one of these half-reactions. The half-reactions involve the creation of charged ions and the capture or release of electrons. Under normal circumstances where the reactants are intimately mixed, these processes occur simultaneously at the same location. However, in a battery the two electrodes are separated by an electrolyte that will allow charged ions to pass from one electrode to the other but will not allow electrons to pass. These can only cross from one electrode to the second to complete the reaction through an external circuit. This is the electrical current that can be used to drive electrical and electronic equipment. It is this separation of processes by the use of a selective filter (in this case, the electrolyte) that allows the cell to generate power.

As mentioned, batteries are generally divided into primary cells and secondary cells. Primary cells contain reactants that will only react once to produce power. After that the cell is spent. A secondary cell can be recharged by applying a reverse polarity to the cell, reversing the chemical cell reaction and regenerating the original cell reactants. It is this type of cell that is of use for energy storage.

There are a number of battery types that can be used for grid and utility applications. The most widely used secondary cell is the lead-acid battery, similar to the type commonly used in vehicles (Figure 10.3). Lead-acid batteries account for roughly half of all secondary cell sales, globally, and are widely available. Another type, nickel-cadmium (NiCad) batteries, are less common. NiCad batteries were used for portable computers until they were superseded by alternatives that were better suited to the duty cycle of such devices. They have also been used for automotive applications, particularly under low-temperature conditions when they behave more reliably than lead-acid cells. Other cell types developed specifically for portable electronic devices include nickel-metal hydride cells and lithium-ion cells. Both are potentially useful for utility storage. A high-temperature battery, the sodium-sulfur battery, has also been popular for utility applications.

In addition to these conventional secondary cells, a variety of devices known as flow cells or flow batteries have also been tested for large-scale energy

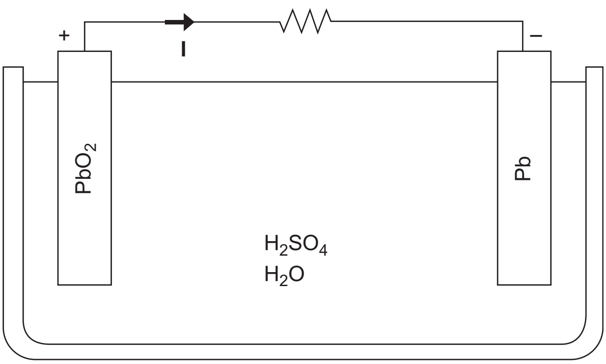

FIGURE 10.3 A lead-acid battery.

storage applications. These include zinc-bromide, vanadium-redox, and polysufide-bromide flow cells. None has so far found widespread commercial application but new types are being developed. They are considered attractive because they have much longer lifetimes than conventional secondary batteries.

Traditional electrochemical storage systems boast a best-case cycle conversion efficiency (electricity to cell storage and back into electricity) of 90%, but a more typical figure would be 70%. Most batteries also suffer from leakage of power over time. Left for too long, the cell discharges itself. This means that battery systems can only be used for relatively short-term storage. Flow cells do not suffer from this problem because the reactants are not stored together and this helps reduce long-term energy losses.

An additional problem with traditional secondary batteries is their tendency to age. After a certain number of cycles, the cell stops holding its charge effectively, or the amount of charge it can hold declines. Much development work has been aimed at extending the lifetime of electrochemical cells but this remains a problem. Again flow cells, because of their design, can avoid this problem.

Against this, batteries are able to supply their output extremely quickly—under 5 msec for a conventional battery and less than 100 msec for a flow battery. Some are also capable of very high-power outputs and discharge rates.

Lead-Acid Batteries

Lead-acid batteries were among the first secondary cells to be developed and were used for load leveling in very early power distribution systems. The cell is based on a reaction between lead-oxide and sulfuric acid. Efficiencies of lead-acid batteries vary depending on factors such as the temperature and duty cycle but are typically between 75% and 85% for DC–DC cycling. However, cells discharge themselves over time so they cannot be used for very long-term power storage. If cycled carefully, cells for utility applications can have lifetimes of 15–30 years.

The cells have a water-based liquid electrolyte and operate at ambient temperature. Both high and low temperatures can reduce their performance. They are also relatively heavy and have a poor energy density, although neither of these factors are a handicap for stationary applications. In addition, they are cheap and easily recycled.

Several very large energy storage facilities based on lead-acid batteries have been built. These include and 8.5 MW unit constructed in West Berlin in 1986 while the city was still divided into East and West, and a 20 MW unit built in Puerto Rico in 1994. While the former operated successfully for several years, cell degradation led to the latter closing after only five years. Lead-acid cells have been very popular for renewable applications such as small wind or solar installations where they are used to store intermittently generated power to make it continuously available.

Nickel-Cadmium Batteries

The nickel-cadmium battery is one of a family of nickel batteries that includes nickel–metal hydride, nickel-iron, and nickel-zinc batteries. There is also a nickel-hydrogen battery in which one cell reactant is gaseous hydrogen. All have a nickel electrode coated with a reactive and spongy nickel hydroxide while the cell electrolyte is almost always potassium hydroxide. Cell reactions vary depending on the second component.

The only nickel-based cell that has been exploited for utility applications is the nickel-cadmium cell. Nickel-cadmium batteries have higher energy densities and are lighter than lead-acid batteries. They also operate better at low temperatures. However, they tend to be more expensive. This type of battery was used widely in portable computers and phones but has now been superseded by lithium-ion batteries.

Efficiencies of nickel-cadmium cells are typically around 70% although some have claimed up to 85%. Lifetimes tend to be rated at around 10–15 years though some have lasted longer. These cells discharge themselves more rapidly than lead-acid cells and can lose 5% of their charge in a month. There can also be a problem with disposal since cadmium is highly toxic.

The largest nickel-cadmium battery ever built is a 40 MW unit in Alaska that was completed in 2003. It occupies a building the size of a football field and comprises 13,760 individual cells.

Lithium Batteries

Lithium batteries, including both lithium-hydride and lithium-ion batteries, have become popular for consumer electronic devices because of their low weight, high energy density, and relatively long lifetimes. Lithium is extremely reactive and can burst into flames if exposed to water, but modern lithium cells use lithium bound chemically so that it cannot react easily. As with nickel, there are a number of lithium cell variants but the most popular today is the lithium-ion cell. These are designed so that there is no free lithium present at any stage during the charging or discharging cycle.

The use of lithium batteries in grid and utility applications is beginning to grow with units being tested in a number of locations. One large installation, due to start operating in 2013, is a 2 MW lithium-ion facility for the Orkney Islands off the northwestern coast of Scotland. The future development of lithium batteries may benefit from interest by automotive manufacturers in their use in hybrid and electric vehicles.

Sodium-Sulfur Batteries

The sodium-sulfur battery is a high-temperature battery (Figure 10.4). It operates at 300 °C and utilizes a solid electrolyte, making it unique among the common secondary cells. One electrode is molten sodium and the other molten

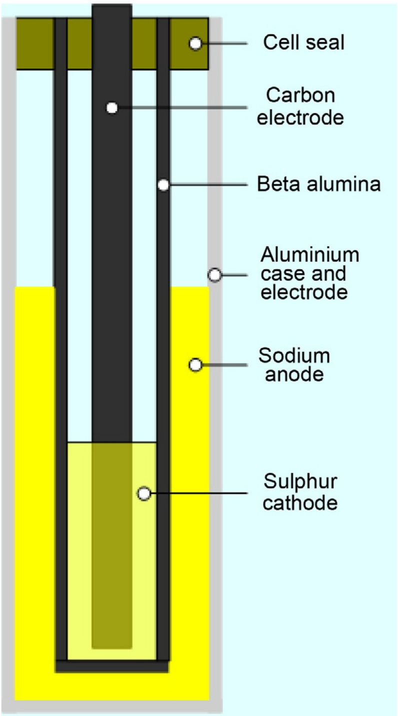

FIGURE 10.4 Sodium-sulfur battery.

sulfur, and it is the reaction between these two that is the basis for the cell reaction. Although the reactants, and particularly sodium, can behave explosively, modern cells are generally reliable. However, a fire was reported in 2012 at a sodium-sulfur battery installation in Japan.

Early work on the sodium-sulfur battery took place at the Ford Motor Co. in the 1960s but modern sodium-sulfur technology was developed in Japan by the Tokyo Electric Power Co., in collaboration with NGK insulators, and it is these two companies that have commercialized the technology. Typical units have a rated power output of 50 kW and 400 kWh. Lifetime is claimed to be 15 years or 4500 cycles and the efficiency is around 85%. Sodium-sulfur batteries have one of the fastest response times, with a startup speed of 1 msec.

Flow Batteries

A flowing-electrolyte battery, or flow battery, is a cross between a conventional battery and a fuel cell. It has electrodes like a conventional battery, where the electrochemical reaction responsible for electricity generation or storage takes place, and an electrolyte. However, the chemical reactants responsible for the electrochemical reaction and the product of that reaction are stored in tanks separate from

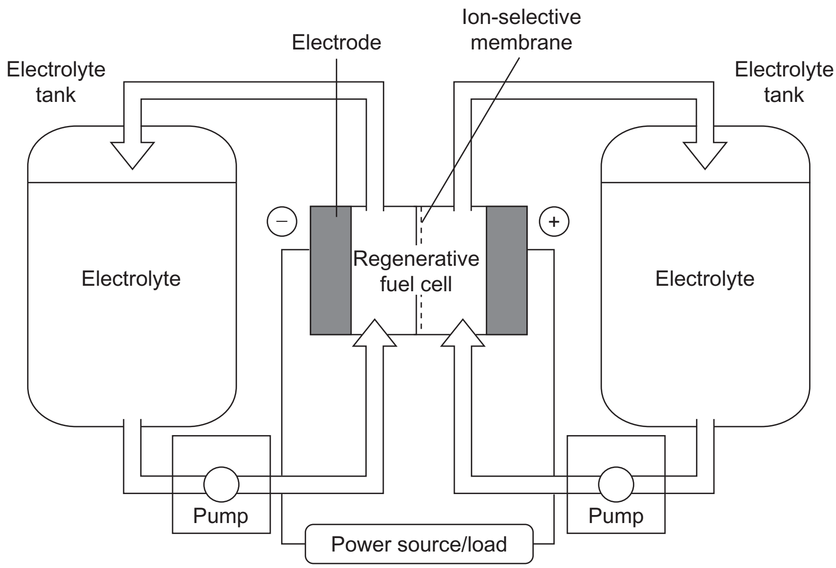

FIGURE 10.5 Diagram of a typical flow battery.

the cell and pumped to and from the electrodes as required, much like a fuel cell (Figure 10.5). The processes occurring here are usually somewhat different to the simple electrochemical process described earlier but the principle is similar.

A number of flow batteries have been tested including the zinc-bromide flow battery, the polysulfide-bromide battery, and the vanadium redox battery. A number of newer designs are also in the research stage. Response times for flow batteries are longer than for conventional secondary batteries but they should be able to supply full power within 100 msec. However, flow batteries have not been extensively tested commercially so their overall performance has yet to be established.

Costs

The costs of battery systems vary widely depending on type and size. Large batteries are generally cheaper than similar small installations. Battery costs are normally based on the cost per kilowatt-hour rather than the cost per kilowatt. The cheapest of the conventional cells is the sodium-sulfur battery with a cost in the range of \$250–900/kWh. Lead-acid battery costs are in the range of \$400–2300/kWh and nickel-cadmium of \$1800–2400/kW. Flow batteries are potentially cheaper with costs in the range of \$150–800/kWh depending on the type.

SUPERCONDUCTING MAGNETIC ENERGY STORAGE

Superconductivity offers, in principle, the ideal way of storing electric power. The storage system comprises a coil of superconducting material that is kept extremely cold. Off-peak electricity is converted to DC and fed into the storage

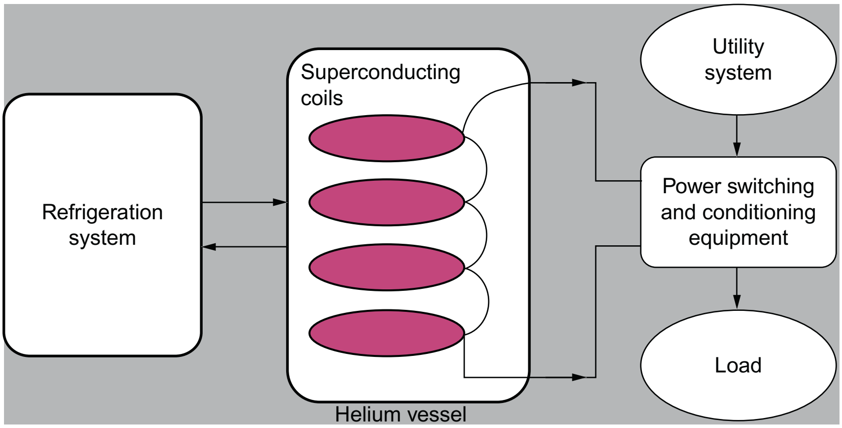

FIGURE 10.6 Grid-connected superconducting magnetic energy storage system.

ring where it stays, ready to be retrieved as required. Provided the system is kept below a certain temperature, electricity stored in the ring will remain there indefinitely without loss (Figure 10.6).

The key to the superconducting magnetic energy storage (SMES) device is a class of materials called superconductors. Superconductors undergo a fundamental change in their physical properties below a certain temperature called the transition temperature, which is a characteristic of each material. When a material is cooled below its transition temperature it becomes superconducting. In this state it has zero electrical resistance. This means that it will conduct a current with zero energy loss.

Unfortunately, the best superconducting materials are metallic alloys that only undergo this transition below 20 °K (−253 °C). Temperatures this low must be maintained by cooling the superconducting coil with liquid hydrogen or liquid helium—in either case an expensive process. In recent years scientists have discovered a new range of ceramic materials that become superconducting at relatively high temperatures, temperatures accessible by cooling with liquid nitrogen. (Liquid nitrogen boils at 98 °K, −175 °C.) Most of these materials have proved to be rather brittle ceramics that are difficult to work with but techniques are being found to exploit them. This is helping make superconductivity more economically attractive for a range of utility applications including storage.

Superconductors store DC current without loss but energy losses occur during the conversion of grid AC to DC and then from DC back to AC. The roundtrip efficiency is typically 90% for daily cycling but will be lower for long-term storage because of the energy required to maintain the coil below its transition temperature. There are also small continuous losses within the coil at the point where power is fed in and out. Startup time for a SMES system is around 5 msec.

When SMES devices were first proposed, they were conceived as massive energy storage rings of up to 1000 MW, similar in capacity to pumped storage hydropower plants. No such storage ring has ever been built, but smaller SMES devices are being used for grid support roles and this appears to be the main commercial market.

Modern technology has enabled smaller commercial SMES storage units with capacities between 100 kW and 100 MW to be constructed. The largest built to date can deliver 10 MW. The storage capacity of these commercial devices is between 10 kWh and 30 kWh, relatively low for utility storage but useful for very fast grid support functions.

One of the earliest SMES devices to be used commercially was commissioned by the U.S. Bonneville Power Administration in the 1980s. This unit had a storage capacity of 30 MJ and a power rating of 10 MW. The device could release 10 MJ of energy in one-third of a second to damp power swings on the Pacific Intertie. Today a typical commercial unit has a storage capacity of 3 MJ (0.83 kWh) and can deliver 3 MW of power for one second.

Costs

Although a few commercial SMES devices are available, costs have been difficult to establish. In general, the cost is relatively low per unit size (MW) but high in terms of storage capacity (MWh). For short-term grid stability duty, they appear to be competitive with other types of storage such as batteries, flywheels, and capacitors.

FLYWHEELS

A flywheel is a simple mechanical energy storage device comprising a large wheel on an axle fitted with frictionless bearings. The flywheel stores kinetic energy as a result of its rotation. The faster it rotates, the more energy it stores. Provided there is a means to extract this energy again, the system can be used for a variety of energy storage applications.

Traditional flywheel-based systems have been in use as mechanical energy storage devices for thousands of years. Millstones, potters wheels, and hand looms have all used them to both store energy and smooth out the peaks in energy delivery by hand or foot. Simple flywheel energy storage devices are also fitted to all piston engines to maintain smooth engine motion. The engine flywheel is attached physically to the engine camshaft, and as the pistons cause the camshaft to rotate they feed energy into the flywheel. For electricity storage applications, energy will normally be fed into the flywheel using a reversible motor-generator.

Flywheel Principle

A flywheel depends on a rotating mass to store energy that is then held in the kinetic energy of rotation of the rotor. The stored energy is proportional to the moment of inertia of the rotor about its axis (directly related to its mass) and its rotational speed.

Conventional flywheels such as those used for piston engines are fabricated from heavy metal disks made of iron or steel. However, these disks are only capable of rotating at low speeds. For power applications, new lighter composite materials have been developed, capable of rotating at 10,000–100,000 rpm without fracturing under the immense centrifugal force they experience. These are made from carbon fiber or glass fiber composite materials that, while less dense than the metals they replace, can store more energy at the very high rotational speeds at which they operate. See Figure 10.7.

Friction losses can be significant at high rotational speeds and energy storage flywheels generally use magnetic bearings to minimize such losses. These are normally supplemented by conventional bearings in case the magnetic bearing fails. Friction losses from air drag can also be significant at the speeds at which modern flywheels rotate, so they will normally be housed in a vacuum chamber, evacuated to reduce this source of loss too. The whole device must be enclosed within an exceedingly strong container that will prevent the pieces of the flywheel scattering like shrapnel in the event of a catastrophic failure.

Energy is fed into a flywheel using a motor to rotate it up to its maximum speed. Energy is extracted from the rotor through a generator driven by the flywheel shaft. Extracting energy from the flywheel will lead to its rotational speed

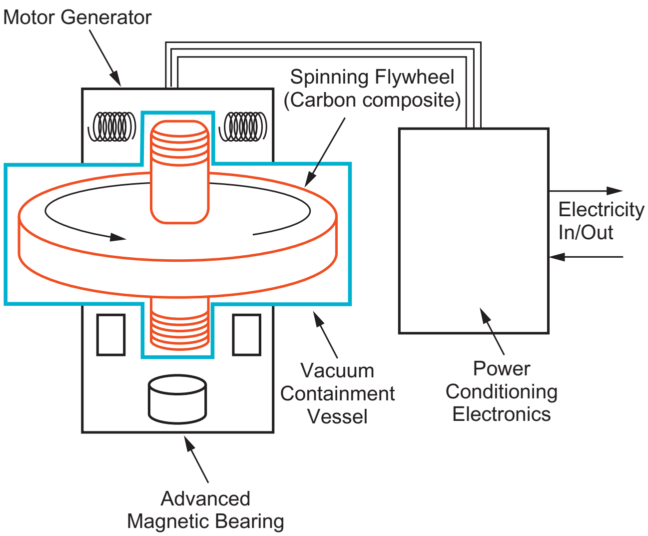

FIGURE 10.7 Flywheel energy storage system.

lowering. The makes grid synchronization virtually impossible so most flywheels use electronic AC–DC–AC converters to eliminate this problem and permit variable-speed operation. Flywheels are generally maintenance free.

Flywheel Performance Characteristics

For a flywheel to operate effectively as an electrical energy storage system it must be kept rotating at its full operating speed. This requires continuous energy feed to compensate for friction losses. Roundtrip efficiency estimates vary with some manufacturers claiming up to 90%. Other sources such as the U.S. Department of Energy suggest 70–80% is more typical.

The energy stored in a flywheel is a function of the mass of its rotor. The larger the rotor, the more it can hold. However, the power it can deliver will depend on the size of the generator used for energy extraction. Most flywheels are designed for backup power or grid support functions during which they are expected to deliver high power for a short period of perhaps a few seconds at most. Such units can have small rotors but relatively large power extraction systems. Other flywheels have been designed for long-term energy delivery. These will have rotors that are capable of storing a large amount of energy compared to the size of the energy extraction system.

Flywheels can usually respond extremely quickly. In grid backup systems they should be capable of reaching full power within half a cycle, 25 msec, and some are quoted with response times of 5 msec. Such units will probably be able to supply their full output between 5 seconds and 15 seconds.

Commercial flywheels are available with power ratings between 2 kW and 2 MW and with storage capacities between 1 kWh and 100 kWh. One of the largest commercial systems is a unit with ten 100 kW flywheels used by the New York Transit System to support its electric traction power network. This system can supply 1 MW of power for 6 seconds. However, the largest flywheel so far constructed is one used in Japan for fusion research. This system can supply 340 MW for 30 seconds.

Costs

The commercial cost of a flywheel system is around $2000/kW, although costs have been put as low as $500/kW. The cost per unit of stored energy is around $1 million/kWh for small commercial flywheel systems though it may be lower for large flywheels. These costs reflect the fact that most flywheel systems store a small amount of energy but can deliver it in a quick burst of high power.

CAPACITORS

Capacitors are electrical or electronic devices that store energy in the form of electrostatic charge. The simplest capacitor comprises two metal plates separated by a small air gap so that no current can pass between them. When a voltage is applied across it the plates become statically charged and this charge can later be released by creating a short circuit between the plates.

Capacitors of various sorts are key components of electrical and electronic circuitry, particularly when creating tuned circuits. While blocking DC current they will allow an AC current to pass with an amplitude that is inversely proportional to the frequency. However, conventional capacitors are capable of storing only a limited quantity of electrical energy. Since the end of the 1970s a new type of capacitor has been available, called an electrochemical capacitor, and these can store much larger quantities of energy. These capacitors, which have names like supercapacitors or ultra-capacitors, are based on electrochemical processes that are similar to those in batteries.

Energy Storage Capacitor Principles

A simple electrostatic capacitor comprises two plates with an air gap between them. When a voltage is applied to the plates, charge builds up on them to neutralize the voltage by creating an equal and opposite static-charge voltage across the plates. The charge will continue to build as the voltage is increased until it is high enough to cause air to breakdown and start conducting electricity.

The amount of charge the capacitor will hold can be increased by placing a dielectric material between the plates that reduces the strength of the field between the plates, allowing more charge to build up before breakdown occurs. An electrochemical capacitor is similar to this in that it has a dielectric material between the capacitor plates, but in this case the dielectric is a liquid electrolyte such as sulfuric acid or potassium hydroxide (Figure 10.8). The capacitor plates themselves are inert materials that will not react with these reagents.

When a charge is applied to the capacitor it causes the usual charge buildup. However, in this case the charge on each plate is neutralized by an opposing layer of charged ions from the electrolyte. This creates a double layer of charge

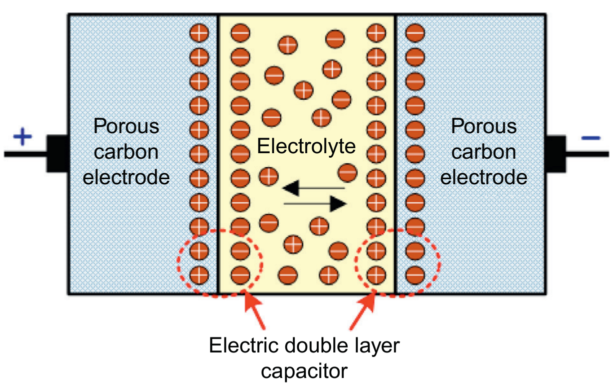

FIGURE 10.8 Schematic of a supercapacitor.

at each plate, effectively leading to two capacitive charged layers, which massively increases the amount of charge the unit can hold.

The simplest electrochemical capacitors use a water-based electrolyte, which restricts the voltage that each can support to around 0.9 V before electrolysis of the water begins. This limits the amount of charge the capacitor can hold but otherwise these capacitors have excellent characteristics. Basic capacitors of this type (known as symmetrical capacitors) have two identical electrodes made from carbon. By varying the construction of one electrode, it is possible to increase the amount of energy the device can hold. Devices of this type are called asymmetrical capacitors. A second type of electrochemical capacitor uses an organic electrolyte, which allows it to support a voltage up to 2.7 V. Both symmetrical and asymmetrical versions of these capacitors are available too. Since single capacitors of either type can only support a relatively low voltage (in grid terms), electrochemical capacitors are usually stacked in series to allow higher voltages to be exploited.

Performance Characteristics

Electrochemical capacitors can be cycled for tens of thousands of times without degradation provided the voltage across them is kept below the maximum so that no internal reaction takes place. However, once they are charged they do lose charge slowly through leakage in the same way as a battery. The leakage levels in water-based electrochemical cells are similar to those of a lead-acid battery. Leakage levels are lower with organic-based electrolytes.

Leakage will reduce long-term storage life, but with fast cycling, roundtrip efficiency can be 95% or higher. In addition, capacitors can normally discharge their energy very rapidly without damage so long as excessive internal heating is avoided. Energy density is moderate between 1 Wh/kg and 5 Wh/kg for symmetrical capacitors and up to 20 Wh/kg for asymmetrical capacitors. In comparison, a lead-acid battery has an energy density up to 42 Wh/kg. Response time is fast at 5 msec.

Units can be designed with the capability to deliver power levels between 1 kW and 5 MW, but actual energy storage capacity is generally relatively low at 1–10 kWh. When power is withdrawn from a capacitor the voltage falls, so sophisticated DC–AC conversion systems are needed to maintain a constant voltage output.

Applications

Electrochemical capacitors have been used both for energy storage and for braking energy recovery systems in automotive applications. For grid use they are best suited to backup or fast reaction grid support, offering a similar performance to flywheels. Although capacitors are not yet widely deployed for grid support they have been tested in a number of configurations. These include adding rapid-response storage to small distributed generation grids or micro-grids where they can provide fast-reacting grid support when the output from intermittent renewable resources suddenly falls and before a backup engine–based system can take over. Capacitors are also being tested for high-voltage grid support services.

Costs

The cost of capacitor storage is likely to be similar to that for flywheels at around $2000/kW. Based on cost per unit of energy storage, the price is again expected to be similar to that of flywheels with costs from $7000/kWh to perhaps $1million/kWh depending on the size of the unit. As with flywheels, this reflects the use of this in low storage capacity, high power output configurations for rapid, short-duration delivery of power.

HYDROGEN ENERGY STORAGE

Hydrogen offers a potential energy storage medium because of its versatility. The gas can be produced by electrolysis of water, making it easy to integrate with electricity generation. Once made, the hydrogen can be burned in thermal power plants to generate electricity again or it can be used as the energy source for fuel cells. In both cases the only combustion product is water. Potentially, it may also be used as an automotive replacement for petroleum or natural gas. Finally, hydrogen has a high-energy density making it an efficient means of storing energy.

For all these reasons hydrogen has been seen as a potential fossil fuel replacement in a future energy economy. For this to become feasible, several problems must be overcome including improving efficiency of its production and finding an economical means of storing it for automotive applications. In the meantime, the limited use of hydrogen as an energy storage medium for intermittent renewable sources such as wind energy is being explored.

Hydrogen Energy Essentials

Hydrogen, as a fuel, can be generated by the electrolysis of water using an electrical voltage. This has been carried out industrially for many years with the main system being an alkaline electrolyzer, exploited most successfully by the Norwegian utility Norsk Hydro. Large-scale electrolyzers have been built that are capable of handling inputs of 100 MW for hydrogen generation and the product is around 99.8% pure. Conversion efficiency is 90%. Alternatives to the alkaline electrolyzer are proton exchange membrane electrolyzers that are currently being developed and could potentially achieve 94% efficiency, but with the need for a platinum catalyst. High-temperature ceramic electrolyzers are also under development.

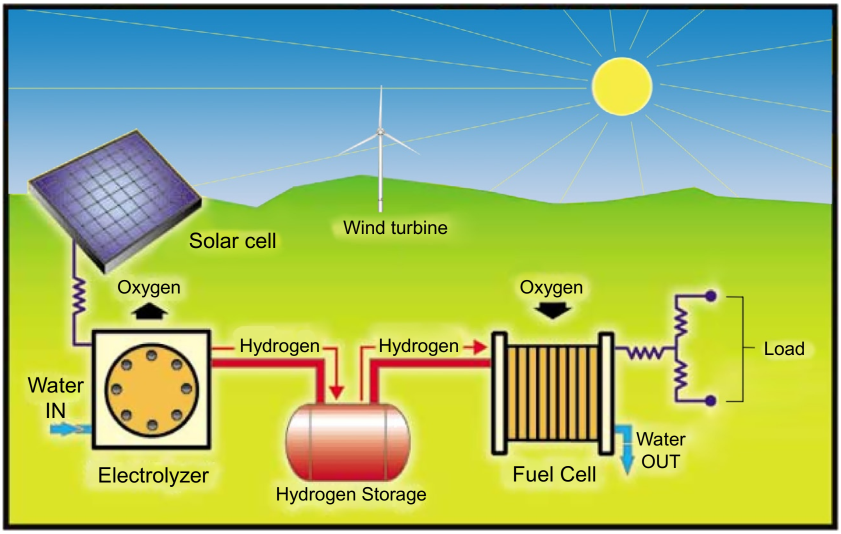

FIGURE 10.9 A hydrogen energy storage system.

Once produced, hydrogen must be stored, for instance, as shown in Figure 10.9. While the gas has a high-mass energy density, it is very light and has a low-volume energy density so it must be compressed or stored in a concentrated state. For power generation applications, storage under pressure in steel or composite tanks is probably the favored method. The gas can be liquefied but only by using cryogenic equipment, making the process costly. There have also been attempts to store hydrogen in the solid state within certain alloys that will absorb it in large quantities. This may eventually offer the best storage method for automotive applications, although the weight of the storage system is currently a major problem.

Once manufactured and stored, hydrogen can be converted back into electricity in a number of ways. It can be burned like natural gas, although the combustion temperature is generally higher than for the former. However, most gas turbines, piston engines, and gas-fired boilers can easily be adapted for its use. When burned in this way it produces only water as a by-product. Alternatively, and possibly most efficiently, hydrogen can provide the fuel for a fuel cell system as illustrated in Figure 10.9. For a future energy economy based on hydrogen, this offers one of the most promising solutions with 60% efficiency achievable in a simple fuel cell and perhaps 70–75% with a hybrid system.

Performance Characteristics

When used as an energy storage medium within an electricity system, hydrogen will generally be used as fuel in a conventional or fuel cell power plant and generally cannot be brought into service as rapidly as some of the fast-acting storage systems discussed above. It should, therefore, be considered as a system of energy arbitrage rather than for grid support.

The main drawback today of hydrogen storage is the round cycle efficiency. With an electrolyzer operating at 90% efficiency and a power plant converting it back into electricity with perhaps 60% efficiency, the best roundtrip efficiency that can be expected is 54%, much lower than other storage systems discussed earlier. Such low efficiency may be tolerable in a renewable energy storage system, such as a wind–hydrogen storage unit where the wind energy must otherwise be shed. It is unlikely to be considered sufficiently efficient for generation from off-peak grid power in most other circumstances if there is an alternative available.

Costs

Since hydrogen energy storage as an electrical energy storage medium has yet to be tested, there are no realistic costs available for practical systems. If it is to be of use it would need to be able to compete with the high storage capacity technologies such as pumped storage hydropower, CAES, or large battery storage.

Chapter 11

Wind Power

Wind power is the second most important renewable source of electric power in the world after hydropower, and since the beginning of the 21st century the total installed capacity has risen rapidly. By the end of 2012 global installed capacity was just under 283,000 MW, around nine times higher than 10 years earlier.[1] During this period wind turbines have developed into a mature power generation technology, while at the same time sophisticated means have evolved to manage their intermittent power delivery into national grid systems.

Most of the new capacity during the past decade has been from onshore wind farms but there is also a growing, and increasingly important, offshore wind sector. Most offshore development has been around European coasts but interest is starting to emerge elsewhere too. Building offshore is more expensive than installing wind farms onshore but this can be balanced by a better wind regime, the ability to build larger wind farms incorporating larger turbines, and the greater ease with which planning consent can be acquired for offshore construction.

Table 11.1 shows the disposition of global installed wind capacity at the end of 2012. As the figures in the table indicate, the greatest concentration of wind capacity was in Europe, the region that has been the strongest supporter of renewable energy during the first decade of the 21st century. As the second decade started, Asian capacity, particularly in China and India, was beginning to rise rapidly. This is likely to become the main area for growth during the second decade. Wind capacity in North America is also high but Africa and South America are notable for their low installed capacities.

The expansion of wind power has resulted in an industry that is global with major manufacturers in Europe, the United States, India, and China. Wind-generated electricity is beginning to prove itself competitive with conventional sources of power generation and can outperform these under circumstances where fossil fuel generation is particularly costly, such as in remote locations. With global competition bringing costs down it is widely predicted that wind power will reach parity with the main conventional sources at some point during the second decade of the century.

1. Wind power figures are from the Global Wind Energy Council.

TABLE 11.1 Regional Disposition of Global Wind Capacity at the End of 2012

Region	Installed Wind Capacity at the End of 2012 (MW)
Africa and Middle East	1135
Asia	97,570
Europe	109,581
Latin America and Caribbean	3505
North America	67,576
Pacific Region	3219

Source: Global Wind Report: Annual Market Update 2012, Global Wind Energy Council, 2013.

WIND RESOURCES

Wind is the movement of air in response to pressure differences within the atmosphere (Figure 11.1). These pressure differences exert a force that causes air masses to move from a region of high pressure to one of low pressure. That movement is wind. The pressure differences are caused primarily by differential heating effects of the sun on Earth's surface, although Earth's rotation will also play a part. Thus, wind energy can be considered to be primarily another form of solar energy.

The effects that lead to the generation of winds are complex and unpredictable, and as a consequence wind is a variable and unpredictable resource. This can make wind power hard to manage on a conventional grid. Advanced

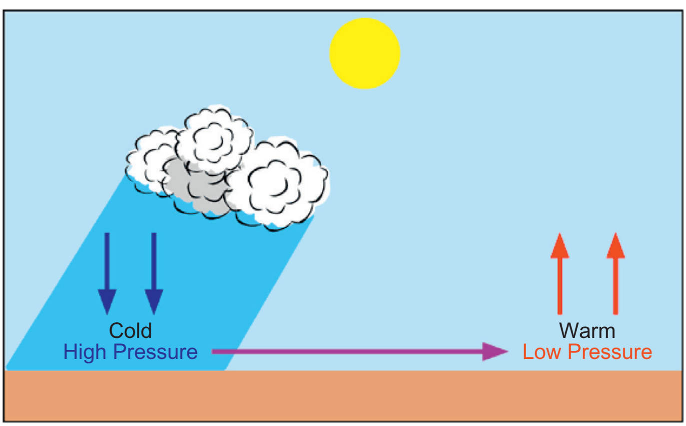

FIGURE 11.1 The generation of wind in the atmosphere.

weather forecasting techniques are rendering short-term wind variations more predictable and this is helping to make the energy generated from wind easier to manage on a grid. Geographical averaging of wind capacity over a specific region can also lead to more reliable output levels, because when the wind does not blow in one part of the region it will often be blowing somewhere else. In this case the larger the region, the more predictable the resource becomes. Nevertheless, wind will always be an intermittent source of energy and this must be taken into account when installing wind capacity.

To assess whether a particular wind site is suitable for exploitation it is necessary to monitor the wind for a period of at least a year, preferably longer. Long-term average wind speeds are usually much more predictable than short-term values, so such long-term assessments provide a much more reliable means of establishing the size of the resource at a particular site.

Annually, over Earth's land masses, around 1.7 million TWh of energy is generated in the form of wind. When oceans are included too, the figure is much higher. Even so, only a small fraction of the available wind energy can be harnessed to generate useful energy. Estimates of how large this might be are difficult; some recent estimates have suggested that global onshore wind could support generating capacities ranging from 100,000 GW to as much as 1,000,000 GW. Even the smaller of these figures is greater than the total global generating capacity from all sources of around 5000 GW. Offshore wind might provide an even greater potential than onshore wind, though exploitation of far-from-shore sites in deep waters is not possible today.

While the potential is clearly large, the exploitation of wind energy is often limited, particularly onshore, by additional restrictions that prevent construction. Urban areas are normally not considered suitable and neither are remote sites if they are close to habitations. Elsewhere, aesthetic considerations may prevent construction of wind farms. In addition, the onshore wind regime will depend critically on local geographic features that can reduce wind speeds in some areas while enhancing them in others. Another factor to take into account is the nature of the ground-cover vegetation, which will affect turbulence levels in the air layers close to Earth's surface. As a consequence, onshore sites need to be chosen carefully and surveyed over an extended period before being developed to ensure they will provide a sufficiently rich wind regime. The best onshore wind sites are often in regions remote from urban centers. This can cause its own problems because these regions are also far from the backbone of the grid system and dedicated transmission capacity may be needed to bring the output to the users.

Offshore construction is usually less constrained than construction onshore, though it should generally avoid shipping lanes or fishing waters. The wind regime offshore is more predictable too, and the relatively smooth surface of the sea means that surface-generated turbulence is often lower than onshore while the extended open spaces allow high wind speeds to develop. Against this, turbines have to cope with regular storms at sea. Like onshore sites, offshore wind farms may also be distant from the grid system and require lengthy transmission connections. However, there are also good offshore sites relatively close to large coastal cities, such as along the eastern seaboard of the United States, and these sites are likely to be attractive for future wind development.

An important factor for consideration at any wind site is that the wind speed increases with height above Earth's surface. In consequence, the taller the tower upon which a wind turbine is mounted, the better the wind regime available. Large wind turbines mounted on tall towers will therefore perform better than smaller turbines on proportionally smaller towers. This wind speed gradient means that the wind speed at the lowest point of the rotor will be smaller than at the highest point. This will create a bending force on the rotor that must be taken into account in wind turbine design.

WIND TURBINE TECHNOLOGY

The modern history of the wind turbine for power generation began during the oil crises of the 1970s. During these early years of wind development many different types of wind turbines were tested. The majority was horizontal-axis wind turbines with a rotor at one end of a shaft and a generator at the other, the whole mounted on the top of a high tower (Figure 11.2). Machines of this

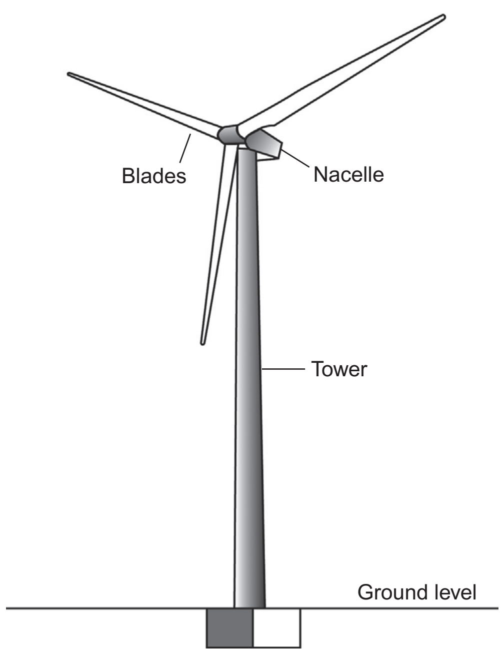

FIGURE 11.2

type were fitted with rotors carrying one, two, three, and more blades. They could be upwind designs, with the rotor facing into the wind and the generator behind, or downwind designs that reversed this arrangement. All used gearboxes to match the rotor speed to generator speed and they often relied on the grid for frequency synchronization and control.

Alongside these horizontal-axis turbines a range of vertical-axis turbines were also developed (Figure 11.3). The most common of these was the Darrieus or eggbeater wind turbine, so-called because its blades were shaped like those of an eggbeater. Other blade designs tested included an H-shaped vertical-axis configuration. The primary advantages claimed for vertical-axis turbines was that they do not need to yaw to keep the rotor facing into the wind while their massive mechanical components—the gearbox and the generator—can be sited on the ground.

In spite of these advantages, vertical-axis machines have never prospered. As the technology has matured most of these designs have disappeared so that today virtually all wind turbines have a similar configuration: a three-blade rotor attached to the front of a horizontal-axis drive-train shaft in an upwind design in which the rotor is always facing the wind. A generator and gearbox (if used) make up the remainder of the drive train that is housed in a protective nacelle mounted on top of a tall steel tower. From an average turbine size of 30 kW in the early 1980s, today's largest onshore machines are in the 2–3 MW range, while offshore machines of 5 MW are now common, and larger machines up to 15 MW are being planned.

Alongside the market for utility wind turbines there is also a parallel market for smaller wind turbines of less than 100 kW. These are often used to supply power to remote sites, for off-grid domestic supply or for a range of small distributed generation applications. One of the largest markets for such small turbines is the United States.

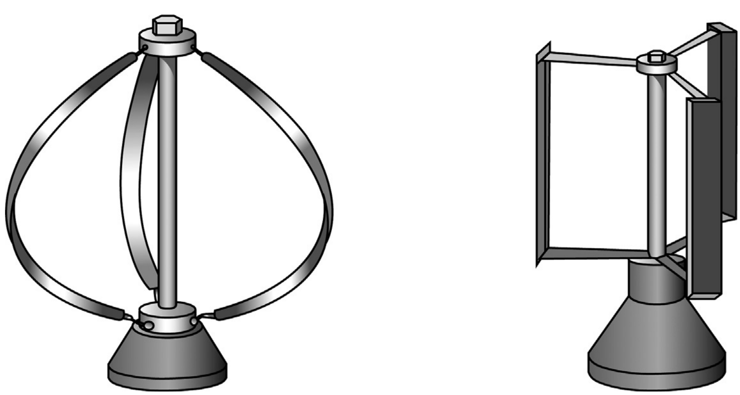

FIGURE 11.3

WIND TURBINE ANATOMY

The standard utility scale wind turbine for both onshore and offshore applications has, as noted already, a three-bladed rotor attached to a drive train and generator with the whole assembly mounted at the top of a tall tower inside a protective housing called a nacelle. The nacelle must be able to rotate, so it is attached to the tower through a yaw bearing that allows the complete structure to turn as the wind direction changes, with the rotor always facing into the wind and the nacelle behind it. In an upwind design, the rotation must be powered by a yaw motor although this is not necessary for a downwind rotor.

The drive train within the nacelle will often include a gearbox that increases the rotational speed of the drive shaft to be able to drive a generator that is synchronized to the local grid frequency, usually 50 Hz or 60 Hz. Other turbines use a variable-speed system with a power electronic converter to ensure the output is always at grid voltage even as the generator speed varies. There are, however, an increasing number of wind turbines that dispense with the gearbox—a component that has often proved unreliable in the past—and use a direct-drive system instead. Direct-drive generators tend to be more expensive but the benefits in terms of higher reliability can outweigh this.

The tower-top structure may have a helicopter pad and will usually be accessible via a ladder or lift within the tower. There may also be a crane fitted to the top for maintenance purposes. Power from the generator will be carried down the tower in cables to a transformer fitted at the bottom of the tower that converts the output to the local distribution grid voltage. If the unit is part of a large wind farm this power may then be carried to a wind farm substation where the voltage is raised further to be fed into the transmission system.

The reliability of modern wind turbines is much higher than it was in the 1980s and 1990s, and most turbine manufacturers aim for lifetimes of 20 years of more. However, anecdotal evidence suggests that maintenance costs rise as machines get older.

Rotors

The rotor is the part of the wind turbine that interacts with the wind and its design will determine the efficiency of the generation unit. The three-blade rotor used on the majority of modern wind turbines represents a balance between cost and efficiency. More blades can, in principle, extract more energy but make the rotor more expensive. Fewer blades are cheaper but lead to balancing problems.

The actual amount of energy that a wind turbine rotor can extract depends on its rotational speed. If the rotor rotates too slowly some wind passes between the blades without energy extraction, whereas if it rotates too fast the turbulence created by one blade will affect energy extraction of the next blade. The optimum rotational speed is usually defined by a parameter called the tip speed ratio

(TSR), which is the ratio of the speed of the blade tips through the air to the wind speed. For a three-blade rotor the optimum TSR is between 6 and 7. It will be clear from this that the optimum rotational speed varies with wind speed, irrespective of turbine size.

For onshore wind turbines the maximum practical unit size is around 3 MW. Beyond this it becomes excessively difficult transporting the massive components to the often remote sites where wind farms are located. Such machines can have rotors up to 120 m in diameter and individual blades up to 60 m. The latter are of a similar size to those used for larger offshore machines up to 5 MW.

Offshore, such 5 MW units are already in use and 10 MW units are being designed. Rotors for these units could have blade lengths up to 75 m, while 15 MW units, now on the drawing board, will require even longer blades. For a given site, onshore or offshore, the selected rotor diameter will also depend on wind speed. A larger rotor will harvest more energy at a low–wind speed site since energy capture will depend on the area swept out by the rotor. In contrast, a smaller rotor can be more economical for a high–wind speed site.

Early turbine blades were often made from wood but most modern wind turbine blades are built from fiberglass-reinforced polymers. Carbon fiber is also being introduced into longer blades to help increase stiffness and strength and this trend is likely to continue. Wind turbine blades are aerodynamically shaped to extract the maximum energy from the wind. The blades must also incorporate features that aid the control of rotor speed. Speed control serves two functions. The first is to enable the optimum rotor speed to be maintained at different wind speeds in variable-speed designs. The second is to ensure that the rotor does not run too fast in high winds. Although turbines would ideally operate in all conditions, if these become too severe, the turbine will normally be shut down completely.

Various methods of speed control are possible. Passive speed control involves designing blades that aerodynamically stall when the wind speed becomes too high, shedding wind. Stalling is a simple technique but it does not help to vary rotor speed with wind speed. The alternative, used by many modern designs, is active pitch control. This involves fitting each blade with a motor at the point where it joins the hub so that it can be rotated about its long axis to change the blade pitch as wind speed varies. Since the optimum rotational speed depends on wind speed, this also allows wind turbines with variable-speed generators to control the speed continuously for optimum efficiency.

Most large wind turbines have a maximum rotational speed of 20 rpm though smaller units may rotate faster. The speed is limited for two reasons. The first is to ensure centrifugal forces do not become too great within the blade. The second is to limit airborne noise, which is a function of blade tip speed. The faster the blade tips move through the air, the more noise they generate. Since blades on large rotors will have a higher tip speed than those on a smaller rotor turning at the same speed, maximum rotational speed on large rotors is normally relatively low.

Although most utility-scale turbine blades adopt a broadly similar shape and structure, there are a number of advanced blade designs under development. These include rotors that have the ability to alter the pitch of each blade independently. This capability may be used to alter blade pitch at different points in the rotational cycle to compensate for the changing wind speed at different heights. Other blades are able to twist under heavy loads, such as during very high gusts, to shed load, a passive system that can help reduce stress fatigue. Other advanced blades are being developed with a number of adjustable sections, each of which is independently controlled by a microprocessor. These complex blade designs also aim to reduce the fatigue loading on blades as well as controlling rotational speed.

Yet another type of new rotor design has the ability to change the length of each blade to create a variable diameter rotor. With a rotor of this type the diameter can be maximized for low wind speed and then reduced as wind speed increases, both controlling rotational speed and reducing the fatigue stress on the rotor blades.

Yawing

The rotor of a horizontal-axis wind turbine must always be oriented so that the plane of rotation is perpendicular to the direction of the wind. This can be accomplished either by having the rotor face the wind with the nacelle behind (an upwind design) or with the nacelle facing the wind and the rotor behind (a downwind design). A downwind design is mechanically simplest because it is possible to use vanes on the nacelle that ensure the orientation is maintained passively simply by the effect of the wind.

Many early wind turbines took advantage of the simplicity of the downwind design but problems with this were soon recognized. The main difficulty arises because of the shadow effect of the tower as each rotor blade passes behind it. This leads to a momentary drop in wind pressure, generating additional fatigue stress in each blade. Noise problems can also arise from the same source. In consequence, modern designs adopted the upwind orientation.

Precise upwind orientation is important to avoid uneven stress on the rotor that can lead to other forms of fatigue. Maintaining an accurate upwind orientation requires that the turbine be equipped with a yawing motor to turn the nacelle. Modern turbines usually use a stepwise system of yawing to keep pace with any changes in wind direction.

The yawing motor also serves a further function. If the nacelle turned continuously in one direction to face the wind the cables from the top of the tower to the bottom would soon become twisted. The yaw motor enables this situation to be avoided by alternating the direction of the yaw as necessary.

Drive Trains and Generator

The drive train of a wind turbine begins with the shaft to which the rotor is attached (Figure 11.4). This transmits the mechanical energy generated by the rotor in the form of a rotational force or torque. In most early wind turbines and in many modern units the shaft is connected to a gear box that increases speed of rotation from perhaps 20 rpm to 1000 rpm or 1500 rpm (50 Hz) or 1200 rpm or 1800 rpm (60 Hz), suitable to drive a synchronized generator. A drive shaft from the gearbox is then linked to the generator.

The drive train has to endure more than simply the rotational torque produced by the rotor. The force of the wind on the rotor blades can be extremely uneven and this will generate lateral or bending forces too, which are transmitted into the gearbox and generator. While shock-absorbing components can help reduce the effect of such lateral forces, the effect on the gearbox can often be severe and this can lead to early failure.

Various attempts have been made to improve gearbox reliability but perhaps the optimum, if most expensive, solution is to remove the gearbox all together and drive the generator directly from the rotor. Direct-drive generators are becoming increasingly popular in large wind turbines and work is being carried out to develop a superconducting direct-drive generator for large offshore wind turbines of 10 MW or more.

The generators used in early wind turbines were asynchronous generators (often motors operated in reverse) that relied on the grid to control their rotational

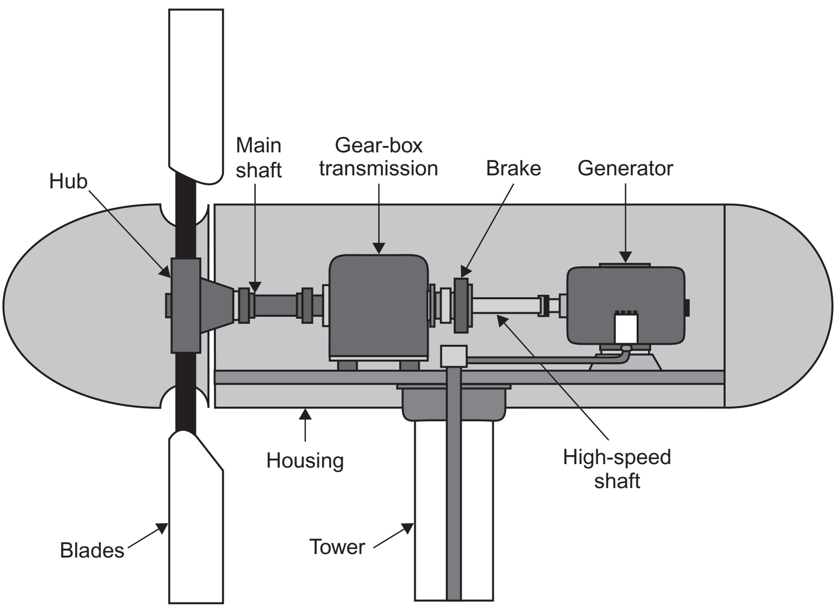

FIGURE 11.4

frequency. This generally weakened the grid and large wind farms with this type of turbine usually required some form of reactive compensation to improve grid stability. Modern wind turbines are normally required to be able to maintain their synchronization with the grid independent of the grid itself, so a more sophisticated design is necessary. This has the added advantage of allowing them to help maintain grid stability rather than reducing it.

Designing an effective generator system for a wind turbine can be difficult because of the variable wind speed conditions. A conventional synchronous generator can only rotate at one speed if it is to supply power at the grid frequency so that the wind turbine must be maintained at a single rotational speed. To overcome this some wind turbine designs include two generators, one for low-speed operation and one for higher-speed operation.

While this is the cheapest variable-speed solution, maintaining fixed rotational speeds creates additional stresses on the rotor and drive train, something that manufacturers are seeking to avoid to improve reliability and lifetimes. The best way to avoid many of the problems associated with variable wind speed is to use a variable-speed generator. The disadvantage of this is that a generator operating at varying speeds will produce an output of variable frequency. Variable-speed operation can, therefore, only be achieved by using some form of power electronic frequency conversion system to maintain grid frequency independent of the frequency from the generator. These electronic systems convert the output from the generator to direct current and then back to alternating current at the grid frequency.

Two types of variable-speed generators have been used in recent years. The first is called a partial conversion system and uses a doubly fed generator to provide limited speed variation. The second is a full conversion generator that is more expensive but also more flexible.

Variable-speed operation reduces the stress on the rotor because the wind turbine can always operate at the optimum speed for the wind conditions. In addition, it means that energy can be harvested over a wider range of wind conditions than is possible with fixed-speed generators. A further advantage is that variable-speed generators with full AC–DC–AC converters can provide grid frequency support facilities, as noted before. This can make them easier to integrate into modern grids.

Towers

The tower of a wind turbine has to be tall enough to lift the rotor and blades so that the blade tips are both clear of the ground and clear of the layer of turbulent air found close to the ground or sea. This will often require a higher tower onshore than offshore for a similar sized rotor because the turbulent air layer is usually thicker onshore. In some cases the rotor may be lifted higher still to gain access to the higher wind speeds found at greater distance from the ground or sea.

Towers for early wind turbines were often made from a lattice steel structure but modern towers are of tubular construction, generally of steel or concrete. Most today are made from tubular steel sections that can be bolted together at the site. Towers are conical in shape, with the base having a larger diameter than the top. Aesthetically the optimum arrangement is considered to be when the tower height is the same as the rotor diameter.

Tower height is also important as the length of the tower is responsible for one of the key structural resonances of a wind turbine. It is critical that this should not be excited by the rotational frequency of the rotor as it could lead to tower failure. This is not normally a problem with onshore wind turbines because the towers are too short, but it can be with offshore turbines mounted on monopole towers with a substantial length below sea level.

Steel towers for large wind turbines are becoming extremely heavy as turbine sizes rise offshore and alternative structures are being sought. One possibility is to construct towers from prefabricated concrete sections. However, concrete does not normally offer the same structural strength as steel. As well as its load-bearing capability, which must be sufficient to support the tower top nacelle and rotor, tower strength is an important issue because the tower is subject to significant bending forces as well as torsional forces generated by the effect of uneven gusting on the rotor. Both must be resisted without significant fatigue stress.

OFFSHORE WIND TURBINE TECHNOLOGY

Offshore wind power started to accelerate toward the end of the first decade of the 21st century and has become an important part of global wind power expansion. By the end of 2012 global capacity had exceeded 5400 MW. Most of this capacity is in Europe but development in both Asia and the United States are anticipated. China has started to build offshore wind farms, as has South Korea, and Japan is exploring its use as a potential replacement for nuclear capacity.

Wind turbines for use offshore are similar to those used onshore, and offshore wind turbines lean heavily on the technology developed for onshore units. However, offshore turbines have to be more rugged than onshore turbines because of the harsher environment. Additionally, offshore wind turbines can be significantly larger than those used onshore. This is important because using larger units can partially outweigh the additional cost of creating a wind turbine foundation offshore. It is the latter that is primarily responsible for making offshore construction more expensive than onshore.

Offshore construction has several advantages over onshore construction. One of the most important is that there are fewer environmental restrictions so that it is often easier to gain permission for development offshore than for construction onshore. The wind regime is generally better offshore too and this means that similarly sized wind farms will generate more power, more reliably.

As a result, offshore wind development, where it is feasible, is likely to become increasingly important over the next 10–20 years.

Against these advantages, offshore development is more expensive than development onshore. As already noted, the most important additional expense is the construction of the foundation for an offshore wind turbine (Figure 11.5). Then, once the turbines have been installed, they are subject to much more severe conditions than onshore so reliability is a key issue. On top of that, offshore maintenance is much more difficult to carry out and therefore also more costly than for a similar unit onshore. All these factors affect the economics of offshore construction.

There are a variety of structures that can be used to establish an offshore wind turbine foundation. The most common of these is a simple monopole tower, similar to the tower used onshore. The foundation for this type of structure is created by using a pile driver to drive a steel monopole into the seabed. This base structure normally terminates with a flange at just above sea level and the tower to support the nacelle and rotor is bolted to this flange. Monopoles can be used in depths up to around 30 m, provided the seabed is suitable for pile driving, but beyond that depth the driving the base becomes more difficult. For long monopoles, resonance effects may also start to come into play and this will limit their applicability.

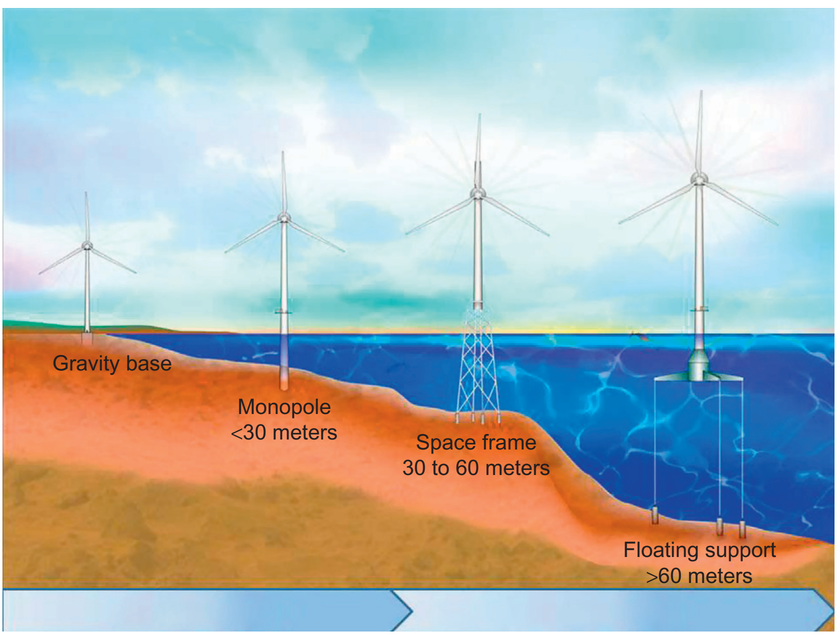

FIGURE 11.5 Offshore wind turbine foundations.

For shallow waters the main alternative to the monopole is the gravity base structure. This is a massive foundation, usually fabricated onshore from concrete sections then floated to the wind turbine site and then sunk to the seabed by loading it with ballast. The wind turbine is then attached to the concrete base that relies on its mass to remain stable. Gravity structures can be used for a variety of seabed conditions but site preparation is necessary to ensure the base sits horizontally. Otherwise, the result is a wind turbine that is not perpendicular.

For deeper waters a range of tripod and space frame foundations can be used. These will have three or more legs, each anchored to the seabed to provide a wide, stable base. The wind turbine is again bolted to the top of the foundation structure. Such structures are the most economical in water depths of 40–60 m.

Stabilizing an offshore wind turbine in water depths greater than 60 m becomes increasingly expensive, and at this depth the only real solution is some form of floating support. No major wind farms have yet been deployed with floating supports but a number of designs are being tested. These include fully buoyant platforms that are anchored in deep water and can support multiple wind turbines; single turbine supports including spar buoys that are partially submerged and anchored to the seabed; and tension leg platforms with several legs that are partially buoyant but are held in position under water by steel guys under tension.

Wind turbines for offshore use have to be more rugged than onshore units so that they can resist the corrosive effects of seawater. Maintenance is much more difficult to carry out offshore too because of the access problems. To improve offshore reliability wind turbine manufacturers are developing remote monitoring and control systems that can both identify and manage faults as they develop so that units can be kept in service and until maintenance can be carried out.

Another issue with offshore development is the means of bringing power ashore. Most of the early offshore wind farms in European waters use simple AC transmission lines to bring power from an offshore substation to a substation on shore. However, as the distance from shore increases, AC transmission becomes less effective because of capacitive losses and it is necessary to switch to high-voltage DC (HVDC) transmission. The crossover between the two systems, economically, is generally considered to be around 100 km. HVDC transmission lines are beginning to be introduced for offshore wind farms in European waters.

For agglomerations of far-from-shore wind farms there are also arguments for setting up dedicated offshore grids that can be linked to onshore grids at more than one point. Typical of such proposals is the North Sea Supergrid, which would link several littoral countries both to multiple offshore wind farms and to one another.

WIND FARMS

A wind farm is a collection of wind turbines that operate as a single power station. Depending on its size, a wind farm will normally have a dedicated substation into which power from all the wind turbines is fed and from which it is carried to the nearest access point to the grid system (Figure 11.6).

Power can be fed either into a local distribution system, or, for the largest wind farms, directly into the transmission grid. As many wind farms are located in regions remote from existing transmission system backbones, wind power development will often involve additional transmission lines. This can add to the expense of a wind project.

When wind turbines are built close together they will often affect one another because, depending on the wind direction, some turbines will be downwind of others and will therefore experience turbulence from the upwind turbines. Correct turbine placement is vital if a wind farm is to extract the maximum energy from the resource.

Onshore wind farms are generally limited in size and this, together with the nature of the terrain, will often make turbine placement less critical. For offshore construction, however, where individual turbines are larger and the number of turbines can be larger too, placement can be critical and can in some cases affect the economics of a project. Energy losses due to poor placement of up to

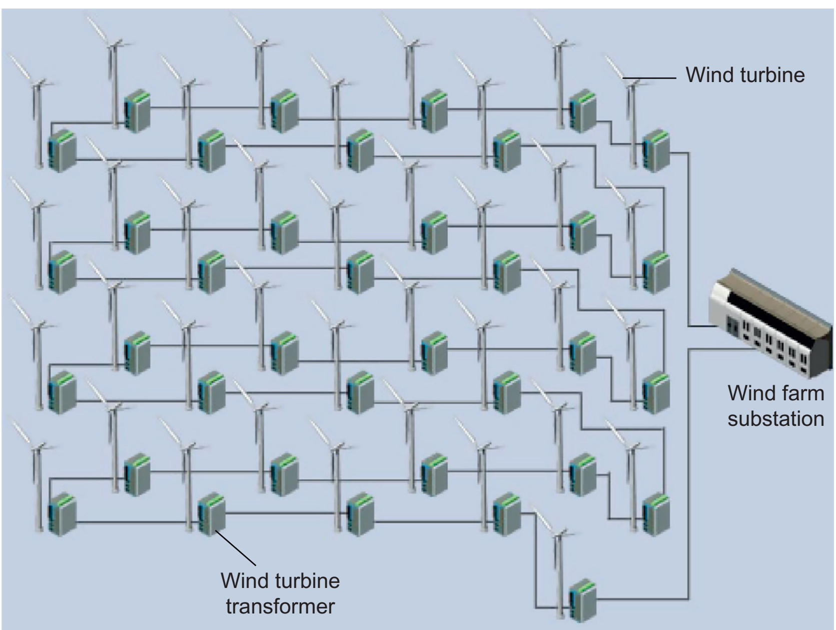

FIGURE 11.6 Layout of a wind farm.

10% have been recorded. Modeling airflow over the wind farm site is the best way to determine the optimum placement, but computer modeling techniques have yet to reach the level of sophistication to accurately model a large offshore wind farm.

Maintaining a good distance between wind turbines means that wind farms can end up occupying a very large area of land or sea. This is not necessarily a handicap because the turbines generally occupy less than 1% of the actual area. The remainder is usually accessible for farming use onshore or for fishing or the passage of ships offshore.

ENVIRONMENTAL EFFECTS OF WIND POWER

The environmental impacts of wind power, with the exception of the aesthetic impact of a series of wind turbines being added to the landscape, are generally limited. Most of the problems that arise when attempting to gain permission to erect a wind turbine or build a wind farm are related to the siting of wind turbines due to the visual impact. This has proved a major issue, particularly in densely populated countries such as the United Kingdom where wind farms are often subject to lengthy permitting procedures. It is partly for this reason that offshore construction is accelerating around European shores.

Of the other potential environmental effects, noise has been considered an issue in the past, but most modern wind turbines with their large, low-speed rotors have limited noise impact provided they are reasonably distant from habitations. Other areas of concern, such as the danger to birds from rotating turbine blades or impacts on marine life from construction offshore, have not generally proved serious and studies have suggested the impact is small, although anecdotal evidence sometimes contradicts this.

WIND INTERMITTENCY AND GRID ISSUES

Electricity generated from wind turbines suffers from two problems that make it difficult to accommodate on a conventional grid. First the wind is an intermittent source of energy so a single wind turbine cannot ever provide electric power continuously. Second, the wind is unpredictable so that it is impossible to know in advance when a wind turbine will supply power and when it will not. Together these make the management of wind energy more difficult than most other sources of power.

Intermittency on its own is not necessarily a problem. Tidal power is intermittent but the output can be predicted with great certainty so that the dispatching of tidal power is relatively easy. Solar power is intermittent too and has a degree of unpredictability, but during daylight hours there will be generally some solar power available so the problem here is less severe. In addition, in regions where the sun shines often, power demand, particularly from air conditioning, will often follow solar output so that solar power can provide a reliable source of peak power that will coincide with rises in this type of demand. Wind power, on the other hand, cannot be relied on to coincide with anything. (However, the wind often blows more strongly in winter than in summer, so wind and solar can be complementary over a seasonal timescale.)

There are various ways of managing the intermittent and unpredictable output from wind power plants. The simplest solution, and the one often applied where wind is used for remote or small domestic supplies, is to store the wind energy—most frequently with a small battery storage system. So long as the wind power–generating capacity is sufficiently large enough compared to the demand, the wind power plant will be able to supply enough stored energy to provide a continuous supply of electricity. Such systems usually include some backup source of electricity too, for those rare periods when there is a lengthy calm period.

Energy storage is, in principle, the most robust solution for wind energy management on grid systems too. Unfortunately, energy storage capacity is generally expensive and very few grids have storage sufficient to manage the large amounts of wind power being introduced into grid systems across the world today. As the proportion of renewable energy grows, it may become economically expedient to expand storage capacity. Meanwhile, novel solutions to this problem for wind have been proposed, including the integration of wind generation with hydrogen production as a means of energy storage. It has yet to be established that this is an economically viable solution.

Where storage is not available, wind output variability is usually managed by maintaining a sufficiently large standby capacity to step in when wind output fails. The cheapest way of achieving this, but one that is not available everywhere, is to use hydropower capacity as backup. Provided this capacity is based on dam and reservoir power plants, hydropower can be brought online or taken offline rapidly and at will. In a sense this is much like pumped storage hydropower but without the full flexibility. Nevertheless, where hydropower capacity of this sort is available, the cost of grid integration of wind energy has been shown to be cheaper than where it is not available. When such hydropower capacity is not available the standby capacity will probably be based on natural gas–fired combined cycle power plants, which are more expensive to operate and therefore increase overall wind integration costs. Coal plant manufacturers are also trying to adapt their technologies to be able to provide the flexibility to provide grid support of this type too.

There are also other means of helping to integrate wind power. While wind unpredictability and intermittency cannot be avoided, there are ways of ameliorating the problems associated with them. One is to use sophisticated weather forecasting techniques. If wind output can be predicted with reasonable accuracy several hours or even a day ahead, dispatching of the power on the grid becomes much easier. Integrating weather forecasting into automated dispatching systems is already being used with advanced dispatching systems and the accuracy and applicability of such techniques can be expected to improve in the future.

A further factor that can help make wind output more predictable is geographical averaging. An individual wind turbine will always have a varying and intermittent output. However, if the output of one turbine is combined with that of a second at a different location, the combined output will generally vary less because the wind level at one location will not be exactly correlated with that in another. This idea can be expanded so that the output from wind farms over a wide geographical area can be considered collectively as one source of power with much less variability than any one wind farm individually can provide. Modern computerized grid management systems can treat groups of power plants such as wind farms as single, virtual power plants, creating more reliable and therefore more valuable wind energy power sources.

WIND CAPACITY LIMITS

It is clear from figures already quoted that there is, in principle, enough wind energy to supply global electricity demand 20 times over, or more. There would appear to be no limit, therefore, to the amount of electricity that might be generated from wind turbines except under exceptional regional circumstances.

What will limit wind power is the amount of wind energy that can be satisfactorily managed on grid systems without endangering grid security and reliability. This limit will depend on the availability of the means of providing energy when the wind does not blow, including energy storage and alternative power sources. Various estimates have been made in the past for the amount of wind energy that can be absorbed, practically, on modern grids. However, the best way of assessing what is practical is to look at the proportion of wind capacity on existing grids where wind penetration is high.

The highest wind penetration in the world in on the Danish grid where 29% of power was provided by wind energy in 2012. Denmark has links with neighboring countries and so can export excess wind power allowing it to use most of the wind energy it generates. Elsewhere, countries have had problems absorbing all the wind energy they generate. Germany, with 11% of its electricity supplied by wind power in 2012, has often had to curtail wind power during periods of high output because its transmission system does not allow large amounts of power to be transmitted from the north, where most wind generation is situated, to demand centers in the south.

Germany's problem can be solved with grid modification. Meanwhile, the International Energy Agency has estimated that a grid such as that in Denmark would be capable of managing around 60% wind penetration—that is, 60% of all electricity could be provided by wind capacity—although others countries might have lower capabilities. This suggests that with appropriate adaptation and grid integration, perhaps as much as 50% of the power on a typical modern grid might feasibly be provided by wind power. Whether that is desirable is probably more a matter of politics than technology.

REPOWERING

As wind turbine technology has evolved, wind turbines have become both larger and more efficient at capturing energy from the wind. As a consequence early wind farms, based on relatively large numbers of small wind turbines, are starting to appear significantly less economical than modern wind farms with smaller numbers of larger wind turbines. Aesthetically the modern wind farms are often more attractive too.

This change is creating a market for the repowering of existing wind farms with new wind turbines. Repowering can be economically viable and it has the attraction of allowing a new, potentially more financially attractive, wind farm to be built at a site where a wind facility already exists, avoiding the need to acquire the various permits that might be needed at a new site.

The earliest repowering took place in California when large numbers of small wind turbines were scrapped during the late 1980s and early 1990s. Repowering has also taken place in Denmark where, before 2002, around 1800 wind turbines were taken down. Meanwhile, Germany began to encourage repowering in 2004 with financial incentives after most of the best onshore sites had already been used.

Typically, repowering of a wind farm with new turbines aims to double the output of the farm while reducing the number of wind turbines by 50%. Repowering also creates a market for secondhand wind turbines that can be refurbished and then sold on. This has led to older turbines from western European countries being re-erected in countries of the Balkans and eastern Europe. Older turbines can also be exported to developing countries for reuse, potentially cutting the cost of introducing the technology to these nations.

COST OF WIND POWER

As with many renewable technologies, wind power is a capital-intensive form of power generation in which most of the investment over the life of the plant is required for its construction, while operational costs are relatively low and there are no fuel costs. The cost of wind turbines fell sharply during the 1990s and at the beginning of the 21st century as technology development and improved manufacturing techniques allowed economies to be made. During the middle of the first decade of this century costs started to level out as the economies made by development were counterbalanced by some steep increases in commodity prices. This led to the prices for the installation of onshore wind generation starting to rise in the latter part of the decade. However, the financial crisis at the end of the decade and the rapidly increasing competition in the wind turbine market led to costs starting to fall again as the second decade began.

The cost of building offshore wind farms is higher than the cost onshore. For U.K. waters the cost is probably close to 50% higher than for building onshore. Added to this, as offshore wind farms have moved into deeper waters farther from shore the cost of installation has risen, offsetting any gains from improvements in technology and installation techniques. Against this, offshore wind technology is at an earlier point in its development cycle than onshore wind technology, and there are likely to be further economic gains to be made in the cost of offshore technology, particularly for foundation construction. Floating wind turbine platforms and supports, in particular, could cut costs significantly if they can be perfected.

The U.K. Department of Energy and Climate (DECC) has estimated that the installed cost of onshore wind turbines in the United Kingdom in 2011 was £1452/kW. For offshore wind turbines, the installed cost in 2011 was estimated to be £2722/kW, as shown in Table 11.2. The U.K. DECC expects the cost of onshore wind turbines to fall only slightly over the next two decades. This view may be modified by the growing global competition between wind turbine manufacturers that could lead to capital costs falling more than had been expected. For offshore wind the U.K. DECC expected a significant fall in costs over the same period so that by 2030 the installed cost was predicted to be around 30% lower than in 2011. Competition for offshore wind turbines is not yet as fierce as for onshore machines but this could change, leading to steeper falls in price than this suggests.

The actual cost of electricity from a wind farm, as measured by the levelized cost, is also shown for U.K. wind farms in 2011. For large offshore wind farms (greater than 5 MW in capacity) the levelized cost was £91/MWh. Offshore wind, at £169/MWh, was significantly more expensive. In the United States, meanwhile, the U.S. Department of Energy found average wholesale wind energy costs for onshore wind capacity in 2012 to be $40/MWh based on new contracts for electricity from wind farms.

TABLE 11.2 Wind Power Costs in the United Kingdom in 2011

Turbine Type	Typical U.K. Capital Cost in 2011 (£/kW)	Levelized U.K. Cost of Wind Power in 2011 (£/MWh)
Onshore wind turbine	1452	91
Offshore wind turbine	2722	169

Source: U.K. Department of Energy and Climate Change.

Based on these figures, electricity from wind power plants remains generally more expensive than power from the best conventional sources such as natural gas–fired combined cycle power plants, existing coal-fired power plants, and established hydropower plants. It is widely expected that this situation could change during this decade and that sometime toward the end of the decade wind power will achieve parity with these other sources. Whether this will happen will depend on a number of unpredictable factors including whether competition within the turbine market continues to lead to falling prices and the impact of shale gas in the United States and elsewhere.

Chapter 12

Geothermal Power

Geothermal energy is the heat contained within Earth's body. The origins of this heat are found in the processes that led to the formation of Earth from the consolidation of stellar gas and dust some 4 billion years ago into a high-temperature ball of matter. Over time the outer regions of that ball cooled but the core remains at a very high temperature. Radioactive decay within Earth's body continually generates additional heat that augments that already present.

The distance from Earth's surface to its core is 6500 km. At the core the temperature may be close to 6000 °C, creating a temperature gradient between the center and the much cooler outer regions. As a consequence, heat flows continuously toward the surface. Most of this heat reaches the surface at a low temperature and cannot be exploited, but in some places a geothermal anomaly creates a region of high temperature close to the surface. In such regions it may be possible to use the energy, either for heating or in some cases to generate electricity.

The geothermal energy that is capable of being exploited at the surface is contained within Earth's solid outer shell, called its crust. Earth's crust is generally around 56 km thick. Starting from the ambient surface temperature, the temperature within the crust increases on average by 17–30 °C for each kilometer below the surface. Based on this, it has been estimated that the top 3 km of the crust contain around 4.3×10^7 EJ of energy, around 10,000 times more than annual global energy consumption. Below the crust is the mantle, a viscous semi-molten rock that has a temperature between 650 °C and 1250 °C. Inside the mantle is the core. Earth's core consists of a liquid outer core and a solid inner core where the highest temperatures are found.

The Earth's crust is not a shell of uniform thickness. Exploitable geothermal temperature anomalies occur where molten magma in the mantle comes closer than normal to the surface. In such regions the temperature gradient within the rock may be 100 °C/km, or more. Sometimes water can travel down through fractured rock to such anomalies and by convective flow carry the heat back to the surface. More dramatically, plumes of magma may rise to within 1–5 km of the surface and at the sites of volcanoes it actually reaches the surface from time to time. However, direct exploitation of this energy source is likely to be difficult. The magma also intrudes into the crust at the boundaries between the tectonic plates that make up Earth's surface. These boundaries can be identified by earthquake regions such as the Pacific basin "ring of fire."

The most obvious surface signs of an exploitable geothermal resource are hot springs and geysers. These have been used by man for at least 10,000 years. Both the Romans and ancient Chinese used hot springs for bathing and therapeutic treatment. Such use continues in several parts of the world, particularly Iceland and Japan. A district heating system based on geothermal heat was inaugurated in Chaude-Aigues, France, in the 14th century; this system is still in existence.

Industrial exploitation of hot springs dates from the discovery of boric acid in spring waters at Larderello, Italy, around 1770. This led to the development of a chemical industry based on the springs. It was here, too, that the first experimental electricity generation from geothermal heat took place in 1904. This led, in 1915, to a 250 kW power plant that exported power to the local region. Exploitation elsewhere had to wait until 1958 when a plant was built at Wairakei in New Zealand and the Geysers development in the United States that began operating in 1960. Global geothermal generating capacity has grown slowly since then. In 2012 there was 11,224 MW of installed geothermal capacity worldwide.

The principal exploiters of geothermal power generation, by country, are shown in Table 12.1. The largest user is the United States with 3187 MW of generating capacity in 2012. The Philippines has 1904 MW and Indonesia 1222 MW. There are also large geothermal capacities in Mexico, Italy, New Zealand, Iceland, Japan, Kenya, and several countries in Central America. According to the Geothermal Energy Association, from which these figures are derived, there are 25 nations exploiting geothermal power. Two that exploited it in the past, Greece and Argentina, no longer do so.

As Table 12.1 suggests, easily accessible geothermal resources suitable for power generation are not widely distributed; neither are they large. Consequently, geothermal power contributes only a small amount to global generation. Even so, geothermal energy is attractive for power generation because it is simple and relatively cheap to exploit. In the simplest case steam can be extracted from a borehole and used directly to drive a steam turbine as the schematic in Figure 12.1 illustrates. Such easily exploited geothermal resources are rare but others can be used with little more complexity. The virtual absence of atmospheric emissions (although geothermal wells can release carbon dioxide) means that geothermal energy is also clean compared to fossil fuel–fired power.

While natural geothermal fields are relatively rare, there is a much larger geothermal potential linked to deep, hot underground rock within the crust. This heat is more expensive to exploit but could potentially offer a far larger generating capacity. Today, however, only experimental exploitation of this resource has been carried out.

GEOTHERMAL RESOURCE

There are three principle categories of geothermal resource. The simplest to exploit is a source of hot underground water (the geothermal reservoir) that either reaches the surface naturally or can be tapped by drilling boreholes.

TABLE 12.1 Exploiters of Geothermal Power Generation

Country	Geothermal Generating Capacity (MW)
United States	3187
Philippines	1904
Indonesia	1222
Mexico	958
Italy	883
New Zealand	768
Iceland	661
Japan	535
Costa Rica	208
El Salvador	204
Kenya	202
Nicaragua	124
Turkey	93
Russia	82
Papua New Guinea	56
Guatemala	52
Portugal	29
China	14
France	16

Source: Geothermal Basics Q&A, Geothermal Energy Association, Sept. 2012.

This is the geothermal source upon which all existing commercial geothermal power plants are based.

Where there are no underground water sources, anomalies in the crust can create regions where the rock close to the surface is much hotter than usual. This hot rock can be accessed by drilling and pumping a heat-transfer fluid into the rock, then bringing it back to the surface. The process has been tested but never exploited commercially.

The third, and potentially the richest source of geothermal energy, is the magma itself. The magma contains by far the greatest amount of heat energy,

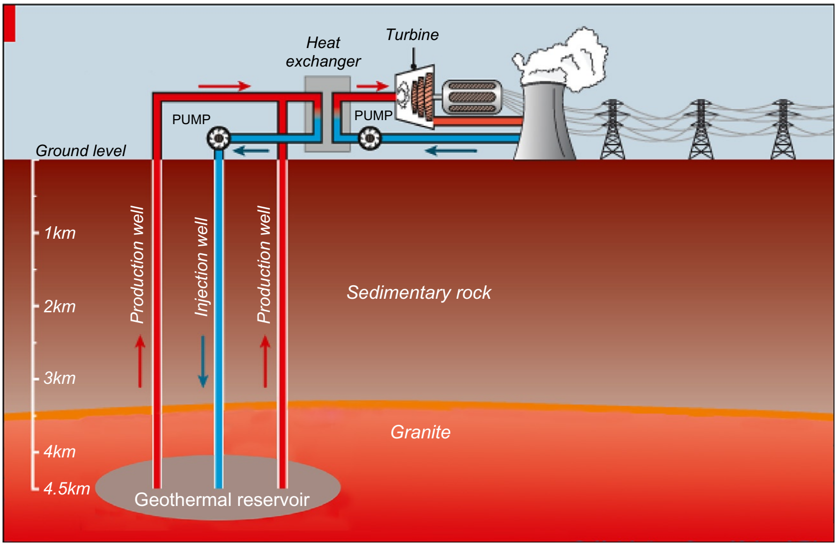

FIGURE 12.1 Schematic of a geothermal power plant.

but because of the temperatures and pressures found within it, this is also the most difficult geothermal energy source to exploit.

Estimating the amount of energy in Earth's crust that could be exploited for power generation is not easy. It has been suggested that there is between 10 and 100 times as much heat energy available for power generation as there is energy recoverable from uranium and thorium in nuclear reactors. Certainly, the resource is enormous, if difficult to access.

GEOTHERMAL FIELDS

Geothermal fields are formed when water from Earth's surface is able to seep through faults and cracks within rock, sometimes to depths of several kilometers, to reach hot regions within the crust. As the water is heated is rises naturally back toward the surface by a process of convection and may appear there again in the form of hot springs, geysers, fumaroles, or hot mud holes. These are particularly common along tectonic plate boundaries.

Sometimes the route of the ascending water is blocked by an impermeable layer of rock. Under these conditions the hot water collects underground within the porous rock beneath the impermeable barrier. This water can reach a much higher temperature than the water that emerges at the surface naturally. Temperatures as high as 350 °C have been found in such reservoirs. This geothermal fluid can be accessed by boring through the impermeable rock. Steam and hot water will then flow upwards through the borehole under pressure and can be used at the surface.

Most of the geothermal fields that are known today have been identified by the presence of hot springs. In the United States, Italy, New Zealand, and many other countries the springs led to prospecting using boreholes drilled deep into the earth to locate the underground reservoirs of hot water and steam that were feeding them. More recently, geological exploration techniques have been used to try and locate underground geothermal fields where no hot springs exist. Sites in Imperial Valley in southern California have been found in this way.

Some geothermal fields produce simply steam, but these are rare. Larderello in Italy and the Geysers in California are the main fields of this type in use today though others exist in Mexico, Indonesia, and Japan. More often the field will produce either a mixture of steam and hot water or hot water alone, often under high pressure. All three can be used to generate electricity.

Deep geothermal reservoirs, as much as 2 km or more below the surface, produce fluid at the highest temperature. Typically, they will produce water with a temperature of 120–350 °C. High-temperature reservoirs of this type are the best for power generation, and the higher the temperature, the more energy can be extracted by a turbine. Shallower reservoirs may be a little as 100 m below the surface. These are cheaper and easier to access but the water they produce is cooler, often less than 150 °C. This can still be used to generate electricity but is more often used for heating.

The fluid emerging from a geothermal reservoir, at a high temperature and usually under high pressure, contains enormous quantities of dissolved minerals such as silica, boric acid, and metallic salts. Quantities of hydrogen sulfide and some carbon dioxide are often present too. The concentrated brine from a geothermal borehole is often corrosive and if allowed to pollute local groundwater sources can become an environmental hazard. This problem can be avoided if the brine is reinjected into the geothermal reservoir after heat has been extracted from it.

Geothermal reservoirs are all of limited extent and contain a finite amount of water and energy. As a consequence, both can become depleted if overexploited. When this happens either the pressure or temperature (or both) of the fluid from the reservoir declines.

In theory, the heat within a subterranean reservoir will continuously be replenished by the heat flow from below. This rate of replenishment may be as high a 1000 MW, though it is usually smaller. In practice, geothermal plants have traditionally extracted the heat faster than it is replenished. Under these circumstances the temperature of the geothermal fluid falls and the practical life of the reservoir is limited.

Reinjection of the brine after it has passed through the power plant helps maintain the fluid in a reservoir. However, reservoirs such as the Geysers in the United States where fluid exiting the boreholes is steam have proved more difficult to maintain since the steam is generally not returned after use. This has led to a marked decline in the quantity of heat from the Geysers. In an attempt to correct this, waste water from local towns has been reinjected into the reservoir. Some improvement has been noted.

Estimates for the practical lifetime of a geothermal reservoir vary. This is partly because it is extremely difficult to gauge the size of the reservoir. While some may become virtually exhausted over the lifetime of a power plant, around 30 years, others appear able to continue to supply energy for 100 years or more. Better understanding of the nature of the reservoirs and improved management will potentially help maintain them for longer in the future.

BRINE-METHANE RESERVOIRS

In some rare cases the hot brine in an underwater reservoir is found to be saturated with methane too. Such reservoirs normally occur in regions rich in fossil fuel. Where such a reservoir is found it is possible in principle to exploit both the heat in the brine and the dissolved methane gas to generate electricity. The only major reservoirs of this type known today are in the Gulf of Mexico.

HOT DRY ROCK

Underground geothermal reservoirs are relatively rare. More normally hot underground rock is not permeated by water and so there is no medium naturally available to bring the heat energy to the surface.

Where hot rock exists close to the surface, it is possible to create a human-made hydrothermal source, as shown in Figure 12.2. This is accomplished by drilling into the rock and then pumping water down through the borehole that has been created. If water is pumped under sufficiently high pressure it will cause the rock to fracture—the process is similar in concept to fracking to release shale gas—creating faults and cracks through which the liquid can flow. (In fact, underground rock often contains natural faults and fractures through which the water will percolate. This often helps with the fracturing.) If a second borehole is drilled adjacent to the first, then water that has become heated as it has percolated through the rock can be extracted and used to generate electricity.

The first attempt at this hot dry–rock technique was carried out by scientists from the Los Alamos laboratory in New Mexico in 1973. Since then experiments have been carried out in Japan, the United Kingdom, Germany, and France. One such project, part of the European Hot Dry Rock Research Project, is at Soulez-sous-Forets in France. Here boreholes have been drilled to 5 km below the surface and temperatures of 201 °C found. The project is now producing a small amount (around 1.5 MW of generating capacity) of electric power.

Based on pilot schemes such as this, estimates suggest that a commercial hot dry–rock system will need to provide 10–100 MW of generating capacity over at least 20 years to be economical. The technology is still in an early stage of development and it is likely to be 10–15 years before commercial exploitation is possible.

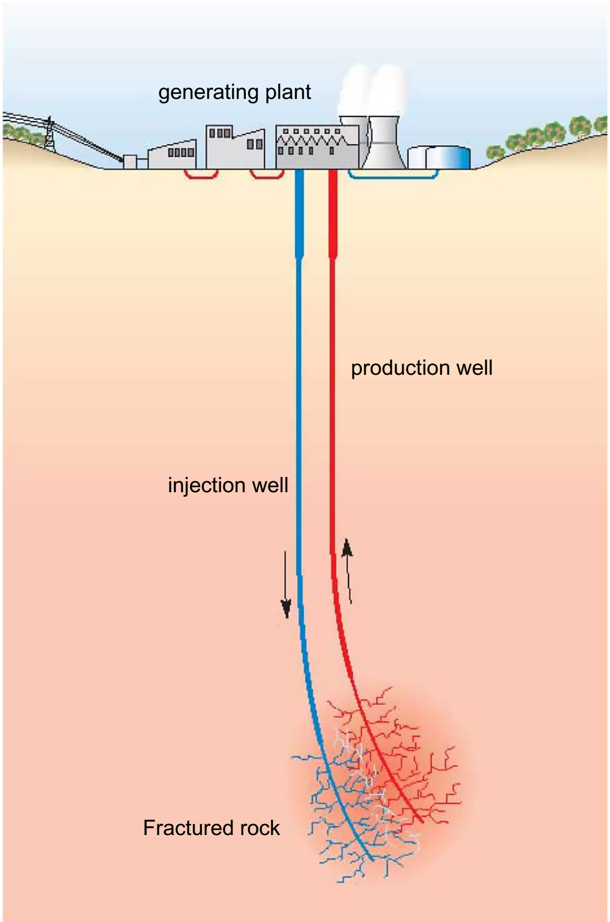

FIGURE 12.2 Schematic of a hot dry–rock power plant.

EXPLOITING THE MAGMA

Extracting energy from accessible magma plumes that have formed within Earth's outer crust is the most difficult way of obtaining geothermal energy, but it is also the most exciting because of the enormous quantities of heat available. A single plume can contain between 100,000 MW-centuries and 300,000 MW-centuries of energy.

Drilling into, or close to, such hot regions is difficult because the equipment can easily fail at the elevated temperatures to which it will be subject. As an additional hazard, if a drill causes a sudden release of pressure, the result can be explosive. Further, ways have yet to be found to tap the heat. Research continues but exploiting magma for power generation is a long-term project with no immediate prospect of exploitation.

LOCATION OF GEOTHERMAL RESOURCES

The easiest geothermal resources to exploit are those that can provide water or steam with a temperature above 200 °C. Resources of this type are located almost exclusively along the boundaries between Earth's crustal plates, in regions where there is significant plate movement. These areas are found around the Pacific Ocean in New Zealand, Japan, Indonesia, the Philippines, the western coasts of North and South America, the central and eastern parts of the Mediterranean, east Africa, the Azores, and Iceland.

Lower-temperature underground reservoirs exist in many other parts of the world and though these contain less energy they can be used to generate electricity too. A project installed in Austria in 2001, for example, generates electricity from 106 °C water, which is also used for district heating. However, these reservoirs can be more difficult to locate in the absence of hot surface springs. Nevertheless, there were around 60 countries using geothermal energy at the beginning of the 21st century for either heating, generating electricity, or both.

SIZE OF THE RESOURCE

Today it is difficult to estimate the size of this energy resource but as survey techniques improve, more accurate data will become available. Based on data available at the beginning of the 21st century, reservoirs located in the United States, for example, might provide 10% of U.S. electricity.

According to figures published by the World Energy Council in its 2010 Survey of Energy Resources the global geothermal generating potential could be between 35 GW and 140 GW, while the technical potential could be 210 GW. However, if hot dry–rock techniques could be exploited, the total potential could be 5–10 times higher than this. A reasonable estimate, again from the World Energy Council, suggests that 8.3% of global electricity generation could be provided by geothermal sources.

GEOTHERMAL ENERGY CONVERSION TECHNOLOGIES

There are three principle ways that have been developed for converting geothermal energy into electricity. Each is designed to exploit a specific type of geothermal resource. The most straightforward of the three is only feasible when a geothermal reservoir produces high-temperature dry steam alone. Under these circumstances it is possible to use a direct steam power plant, which is analogous to the power train of a steam turbine power station, with the boiler of the plant replaced by the geothermal steam source. Provided the steam exiting the reservoir is of suitable quality, this provides an extremely cheap and effective means of generating electricity.

Most high-temperature geothermal fields produce a mixture of steam and hot brine. This mixture cannot be utilized quite so simply and is most effectively exploited using a configuration called a flash-steam geothermal plant. The flash process converts part of the hot, high-pressure liquid to steam and this steam, together with any steam extracted directly from the borehole, is used to drive a steam turbine.

Where the geothermal resource is of a relatively low temperature a third system called a binary plant is more appropriate. This uses the lower-temperature geothermal fluid to vaporize a second low–boiling point fluid contained in a separate, closed system. The vapor then drives a turbine that turns a generator to produce electricity. Although overall efficiency of such binary systems is low, the availability of a cheap heat source makes the system economical to operate.

DIRECT STEAM POWER PLANTS

Dry steam geothermal reservoirs are extremely rare. Where they exist the steam, with a temperature of 180–350 °C, can be extracted from the reservoir through a borehole and fed directly into a steam turbine (Figure 12.3). Steam from several wells of this type will normally be fed to a single turbine to allow an economically large turbine to be used. The pipes that carry the steam from the well heads to the turbine contain various filters to remove particles of rock and any steam that condenses en route. This is the configuration that was used in the very first geothermal power plant in Larderello, Italy, in 1904.

The steam turbine in a direct steam geothermal plant is usually a standard reaction turbine. Unit size in modern plants is typically between 20 MW and

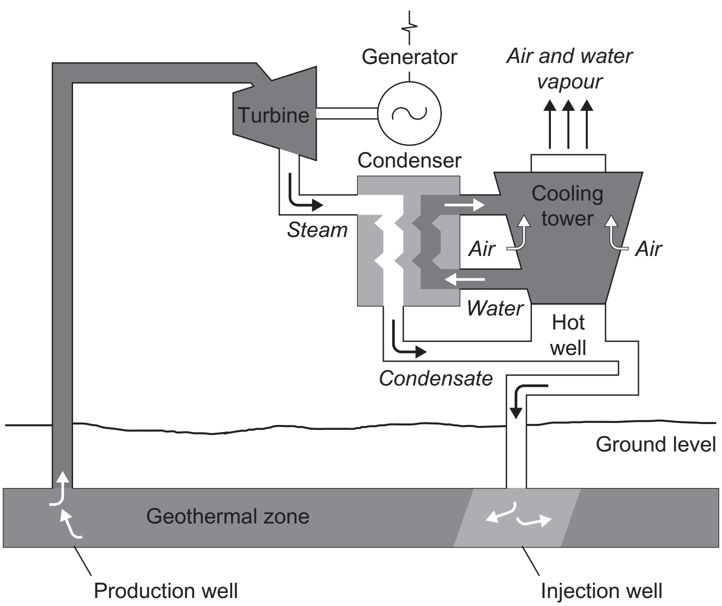

FIGURE 12.3 A direct-steam geothermal power plant.

120 MW. As a result of the relatively low steam temperature and the small size of the turbines, efficiency is generally low at around 30%.

In some cases the steam exiting the turbine may be released directly into the atmosphere. However, the steam usually contains between 2% and 10% of other gases such as carbon dioxide and hydrogen sulfide. Under these circumstances the exhaust from the steam turbine must be condensed to remove the water and then treated to remove any pollutants such as hydrogen sulfide before release into the atmosphere. At the Geysers plant in the United States, the hydrogen sulfide is converted into sulfur as a by-product of the treatment and this sulfur is sold. Condensing the steam also extracts more energy, thereby increasing plant efficiency.

Ideally the geothermal fluid should be returned to the underground reservoir but it is often more economical to release the gas and dispose of the water produced as a result of condensing the steam from the turbine at the surface. Carbon dioxide emissions from such plants could become an issue in the future and, depending on the amounts involved, it may be feasible to capture and either reinject or sequester the carbon dioxide generated from such a plant. In general, however, carbon dioxide emissions are low compared to fossil fuel sources. In the United States, for example, typical emissions from a geothermal plant are equivalent to 91 gCO_2/kWh. For fossil fuel resources the equivalent range of emissions is between 600 gCO_2/kWh and 1000 gCO_2/kWh according to the World Energy Council. Continual removal of underground fluid without replenishment eventually leads to a depletion in the quality of fluid available from the reservoir. At the Geysers geothermal field in southern California, urban waste water has been pumped into the underground reservoir in an attempt to maintain and eventually boost output from the resource.

FLASH STEAM PLANTS

Most high-temperature geothermal reservoirs yield a fluid that is a mixture of steam and liquid brine, both under high pressure (typically up to 10 atm). The steam content, by weight, is typically between 10% and 50%. The simplest method to exploit such a resource is to separate the steam from the liquid and use the steam alone to drive a steam turbine. This would be equivalent to the dry steam plant described earlier. However, using only the steam throws away much of the available energy, particularly where the proportion of steam in the fluid is small.

A more productive alternative is to pass the combined fluid through a valve into a vessel maintained at a lower pressure than the geothermal fluid from the reservoir (Figure 12.4). The sudden reduction in pressure "flashes" a proportion of the hot liquid to steam, creating a larger quantity of steam than was available previously. All this steam can then be separated from the liquid and used to drive a steam turbine. The steam exiting the steam turbine must be treated in exactly the same way as the exhaust from a direct steam geothermal plant to

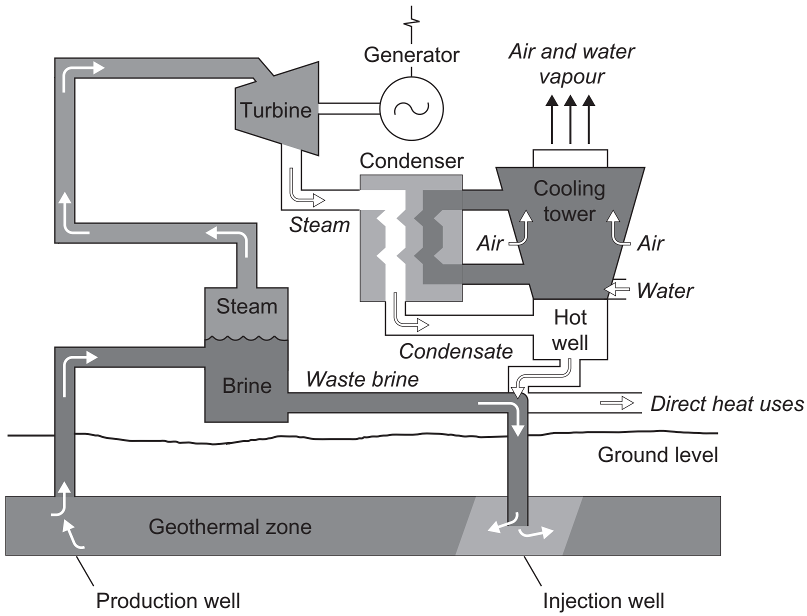

FIGURE 12.4 A flash-steam geothermal power plant.

prevent atmospheric pollution. The remaining liquid, which contains high levels of dissolved salts and presents a potential pollution problem, is usually injected back into the geothermal reservoir.

A further refinement to the flash steam plant is called double-flash technology. This involves taking the fluid remaining after the first flash process and releasing it into a second vessel at even lower pressure. Double flashing results in the production of more steam that can be fed to a second, low-pressure turbine or injected into a later stage of the turbine powered by the steam from the first flash. A double-flash plant can product up to 25% more power than a single-flash plant. However, it is more expensive and may not always be cost effective.

Flash technology plants will generally return a much higher percentage of the geothermal fluid—up to 85% for a single-flash plant and somewhat less for a double-flash plant—to the geothermal reservoir. This will include most of the dissolved chemicals contained in the original fluid. However, some reservoir depletion will still take place and without action this is likely to lead to a falloff in output from the reservoir with time. Capacities for flash geothermal power plants are normally between 20 MW and 55 MW.

BINARY CYCLE POWER PLANTS

Direct steam and flash geothermal power plants utilize geothermal fluid with a temperature between 180 °C to 350 °C. If the fluid is cooler than this, conventional steam technology will normally prove too inefficient to be economically viable. However, energy can still be extracted from the fluid to generate power using a binary cycle power plant. This type of plant has two fluid cycles: the first involves the low-temperature fluid from the geothermal resource and the second is a turbine cycle involving a low–boiling– point fluid (Figure 12.5).

In a binary power plant the geothermal fluid is extracted from the reservoir and immediately passed through a heat exchanger where the heat it contains is used to volatilize a secondary fluid. This secondary fluid is contained within a second, completely closed cycle system. The fluid may be an organic liquid that vaporizes at a relatively low temperature or, in the case of the Kalina cycle,[1] a mixture of water and ammonia.

The vaporized secondary fluid is used to drive a small turbine from which power can be extracted with a generator. Once it has exited the turbine the secondary vapor is condensed and then pumped through the heat exchanger once more. Thus, the cycle is repeated. The geothermal fluid exiting the heat exchanger is, meanwhile, reinjected into the geothermal reservoir. Since 100% of the fluid is returned underground, this type of geothermal power plant has the smallest environmental impact.

Typical binary plant unit size is 1–3 MW, much smaller than for the other types of geothermal technology. However, the small modular units often lend

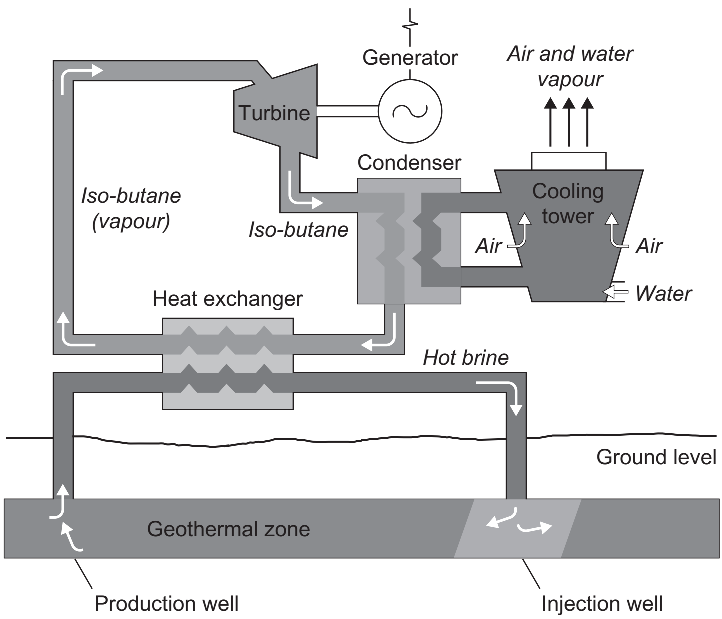

FIGURE 12.5 A binary geothermal power plant.

1. The Kalina cycle is a special thermodynamic cycle designed to obtain maximum efficiency from low-energy resources such as low-temperature geothermal fluids.

themselves to standardization, reducing production costs. Several units can be placed in parallel to provide a plant with a larger power output.

Although the normal application for binary technology is to exploit a low-temperature geothermal resource, the technology can also be used to generate more power from a flash plant. In this case the fluid left after flashing is passed through a heat exchanger before reinjection, allowing extra energy to be taken to power a small binary unit. Adding a binary unit to a conventional flash plant increases the cost but the resultant hybrid plant will have a larger power output.

GEOTHERMAL POWER AND DISTRICT HEATING

While high-temperature geothermal resources are suitable for power generation and lower-temperature resources are often exploited for district heating, it is possible to combine the two. This is typical, for example, of geothermal power plants in Iceland where there is high demand for both power and heat.

Iceland's geothermal reservoirs generally provide a hot brine and energy is extracted from this with a flash plant as described previously. However, after the brine exits the flash plant it still contains a significant amount of heat energy that can be exploited to provide heat for a district heating system. For this, the brine is passed through a series of heat exchangers through the second side of which fresh water is passed and heated. The hot water is then stored in a large tank from which it can be taken to supply the district heating system. The brine, upon exiting these heat exchangers, is reinjected into the reservoir.

FINDING AND EXPLOITING GEOTHERMAL SOURCES

The exploitation of geothermal power is often considered to be more risky than development of other renewable resources because of the uncertainty and high costs associated with finding suitable geothermal fields. Initial surveys of geothermal resources are often carried out by national institutions, but for developers, these then need to be backed up with test wells to determine the exact nature of the resource available. Data from Pacific Rim volcanic regions suggests the presence of a single hot spring will provide a 50% change of an exploitable geothermal field. A boiling spring or fumarole increases the probability to 70%.[2]

Having identified a suitable surface site, prefeasibility studies are likely to cost around $1 million, with a 30% change of failure. Test drilling, usually three wells at up to $2 million per well, has a similar prospect of failure. This risk can be reduced by careful surface study followed by prioritization of the available sites. Such an approach has led to success rates for well drilling in excess of 83% in countries such as Indonesia, Kenya, and New Zealand that have

2. World Bank geothermal assessment.

high-temperature resources. However, success rates can be much lower where low-temperature resources are concerned.

Once a usable underground reservoir has been located, its size must be determined. This involves fluid withdrawal over a long period; indeed, it may not be until several years after production has started that a good picture of the resource can be obtained. Careful sizing of the geothermal plant to match the reservoir size will prolong the lifetime of a reservoir. However, this may not be possible if the plant has to be constructed before full data is available and, as noted, the data may not be available anyway until production has started. Oversized plants such as that installed at the Geysers in the United States have led to a premature fall in output, which will have an impact on overall economics.

COST OF GEOTHERMAL POWER

In common with many renewable resources, geothermal power generation involves a high capital outlay to establish the facility but extremely low fuel and operating costs. In the case of a geothermal plant there are three initial areas of outlay, prospecting and exploration for the geothermal resource, development of the steam field, and the cost of the power plant itself.

As already noted, the cost of identifying a suitable geothermal reservoir is likely to be several million dollars and may not always be successful. If the resource that is found then turns out to be small this will create a much heavier burden on the project than if the resource is large. Steam field development will depend on plant size and will generally be priced out according to the number of boreholes that are to be drilled. For example, the Hellisheidi geothermal power plant in Iceland, with a generating capacity of 300 MW, has 50 boreholes.

The capital cost of building a power plant to exploit a geothermal resource will depend on the quality of the resource. A good resource will have a temperature above 250 °C and good permeability of the reservoir, so it will provide a good fluid flow. Ideally, it will provide either dry steam or steam and brine, the latter being relatively noncorrosive and with low gas content. For good-quality resources of this type today the lowest cost in the United States is around $2400/kW according to the U.S. Energy Information Administration (EIA). Similar costs in Europe are €3000–4500/kW. The lowest cost will be for a plant of more than 30 MW with smaller plants being relatively more expensive to build.

A poor-quality resource will have a temperature below 150 °C, or it could provide fluid at a higher temperature but with some other defect such as corrosive brine or poor fluid flow. Resources of this type will generally only be capable of supplying heat for a small geothermal plant of less than 5 MW and the capital cost could be up to twice that of a large facility based on a good resource.

Further indirect costs will be incurred, depending on the location and ease of access of the site. These will vary from 5% for an easily accessible site and a local skilled workforce to 60% of the direct cost in remote regions where skilled labor is scarce. All these costs will be part of the initial investment required to construct a plant.

The cost of electricity from the geothermal plant will depend primarily on the capital cost of building the plant and the cost of financing the construction. There is also the potential for an additional cost resulting from a local rent for exploiting the resource but that is rare. European estimates put the levelized cost of geothermal electricity at €40–100/MWh. In the United States the EIA has estimated that the cost of energy from a new geothermal plant entering service in 2017 is $100/MWh. These levelized costs make geothermal power competitive with virtually all other forms of power generation.

Chapter 13

Solar Power

Solar energy is the most important source of energy available to the Earth and its inhabitants. Without it there would be no life. It is the energy that drives the photosynthesis reaction. As such, it is responsible for all the biomass on Earth's surface, which is the source of fossil fuels, the products of photosynthesis millions of years ago that became buried beneath Earth's surface. Solar energy creates the world's winds and is, therefore, responsible for wind power; it evaporates the water that is responsible for rain and, therefore, the source of hydropower; even waves and ocean thermal power are both a result of insolation. In fact, apart from nuclear energy, geothermal energy, and tidal power, the sun is wholly or partly responsible for all the forms of energy that are exploited by humans.

While all these different sources of energy, each derived from the sun, can be used to generate electricity, solar energy itself can be used to generate electricity too. This can be achieved most simply by exploiting the heat contained in the sun's radiation. The use of solar thermal energy has a long history.

Solar thermal power exploitation—that is, the use of the sun as a source of heat—can be traced back at least to Archimedes, but its application as a means to generate power is more recent. Among the early examples were attempts made during the 19th century to use parabolic reflectors to concentrate the sun's heat energy and raise steam for a steam engine. At the start of the 20th century solar energy was harnessed to drive an engine that pumped irrigation water for agricultural use in Egypt while the first solar thermal power-generating plant was built in Italy in the 1960s. However, the real driving force for the development of solar thermal power generation was the energy crises of the 1970s.

Electricity can also be generated directly from sunlight using an electronic device called the photovoltaic or solar cell. This second route for converting sunlight into electricity can also be traced to the 19th century to the work of the French scientist Antoine-César Becquerel. He was the first to observe the photovoltaic effect during which a voltage is generated when light falls upon an electrode. Following this, the first true solar cell was built at the end of the 19th century by Charles Fritts who coated gold onto selenium to capture light energy. Fritts' device was very inefficient and it was not until the development of the silicon solar cell by Russell Ohl in 1941 that efficient solar energy collection by this means appeared feasible. This was demonstrated in 1954 when three scientists at Bell Labs in the United States—Gerald Pearson, Calvin Fuller, and Daryl Chapin—produced a cell with an energy conversion efficiency of 6%, and not long afterwards solar cells were adopted as power sources in the U.S. space program.

Today the solar photovoltaic cell or solar cell is rapidly becoming one of the most important sources of renewably generated electricity, and its global installed capacity is third among renewable technologies after hydropower and wind power. Solar thermal power generation has developed less rapidly but it too is beginning to show stronger growth and has the potential to become a major source of power generation later in the century.

SOLAR ENERGY RESOURCE

Solar energy is generated by nuclear reactions within the body of the sun. This energy reaches Earth's surface in the form of electromagnetic radiation. The composition of this radiation as it travels through space toward Earth is around 56% infrared, 36% visible radiation, and 7% ultraviolet, with the remainder belonging to regions of the electromagnetic spectrum outside the energy ranges covered by these three.

Not all this radiation reaches Earth's surface. Some is scattered by dust and molecules in the atmosphere. This scattering is a random process, sending the radiation in all directions so that much goes directly back into space. The remainder reaches the surface, but as diffuse, indirect radiation. Clouds act to reflect more sunlight back into space and they play an important role in regulating the temperature of Earth's surface.

Another part of the radiation is absorbed by molecules such as water, carbon dioxide, ozone, and oxygen within the atmosphere. Water and carbon dioxide absorb energy from the infrared region, whereas oxygen and ozone absorb from the ultraviolet. All these interactions reduce the solar energy flux by around 40% while at the same time changing its composition so that the sunlight that reaches Earth's surface comprises 50% visible radiation and 47% infrared.

The amount of energy carried by solar radiation is normally expressed in terms of the solar constant that measures the quantity of solar energy passing through one square meter of space perpendicular to the direction of travel of the radiation at the average distance of Earth from the sun. According to the World Energy Council, the value of this constant is 1367 W/m^2. When absorption and scattering is taken into account, the total solar flux reaching Earth's surface is estimated to be 1.08×10^8 GW, and the total reaching Earth's surface each year is 3,400,000 EJ. This is between 7000 and 8000 times annual global primary energy consumption. If 0.1% of this energy was converted into electricity with 10% efficiency, it would provide four times more generating capacity than the global total of around 5000 GW. See Table 13.1.

To put this into a more practical perspective, take the example of a group of solar thermal power plants that were built in California in the late 1980s and

TABLE 13.1 Solar Energy and the Earth

Solar constant at the distance of Earth from the sun	1367 W/m^2
Total solar energy reaching Earth in a year	3,400,000 EJ
Total solar flux reaching Earth	1.08×10^8 GW
Average solar energy density at Earth's surface	170 W/m^2

Source: World Energy Council.

early 1990s. These plants were designed on the basis of a solar input of 2725 $kWh/m^2/y$, or 22.75 GWh/y for each hectare. Assuming this energy can be converted with 10% efficiency into electricity, 10 million hectares (100,000 km^2) would be sufficient to supply the entire United States.

The solar input is a key parameter when planning a solar thermal power plant. This is determined by the thermal energy density at Earth's surface, a factor that varies with position on Earth. The average (which takes account of the fact that any point on Earth's surface only receives sunlight for about half the time, the other half being hours of darkness) is 170 W/m^2, while the largest, near the Red Sea, is 300 W/m^2. (The latter is just under one-quarter the value of the solar constant.) Over a year, the average incident energy is equivalent to a barrel of oil or 200 kg of coal.

The sunlight that reaches Earth's surface is of two types: direct radiation and diffuse radiation. The latter is the result of various scattering and absorption processes that take place as the sunlight passes through the atmosphere. Vegetation is able to absorb both types of radiation, and so can solar cells. However, a solar thermal plant requires direct radiation to operate effectively. This limits the applicability of this technology to regions where there is low average annual cloud cover. With no cloud cover, and whatever the location, between 80% and 90% of the sunlight reaching the surface will be direct radiation.

SOLAR SITES AND LAND RESOURCES

Since solar energy conversion depends on solar intensity, solar plants are best sited in regions where there is a good level of insolation. For solar thermal plants, the regions of the globe where cloud cover is rare will offer the best locations. These are often the arid, desert regions of the world where land may have little other use. As a consequence of their nature, these regions are often distant from the major demand centers so transmission system extensions are likely to be necessary to exploit them. Some of the areas with good solar thermal potential include the southwestern United States, Rajisthan in India, the Middle East, North Africa, and the southern-most states of the European Union such as Italy, Spain, and Portugal. Typical space requirements for solar thermal plants are 2–5 ha/MW.

Solar photovoltaic power plants are less sensitive to their situation since they do not require direct sunlight and can operate effectively even with cloud cover. As a consequence, they can be used in a wider range of locations. Large, utility-scale, solar cell arrays capable of generating tens or even hundreds of megawatts of power are becoming popular, but most of the global capacity today is derived from rooftop-mounted solar cells. These units often supply power locally and collectively they represent one of the largest forms of distributed generation in operation. Space demands for both rooftop and utility arrays of solar cells are similar, but rooftop deployment allows cells to exploit otherwise unused space. Utility arrays demand a similar area of land to solar thermal power plants.

SOLAR POWER GENERATION TECHNOLOGIES

As already outlined, there are two ways of turning the energy contained in sunlight into electricity. The first, called solar thermal power generation, involves using the sun simply as a source of heat. This heat is captured, concentrated, and used to drive a heat engine. The heat engine may be a conventional steam turbine, in which case the heat will be used to generate steam, but it could also be a closed-cycle turbine system using an organic thermodynamic fluid, a gas turbine, or a Sterling engine.

The second way of capturing solar energy and converting it into electricity involves use of the photovoltaic or solar cell. The solar cell is a solid-state device like a transistor or microchip. It uses the physical characteristics of a semiconductor such as silicon to turn the sunlight directly into electricity. The simplicity and durability of the solar cell makes it an extremely attractive method of generating electrical power.

As with several other renewable technologies, solar energy is intermittent; it is only available during hours of daylight. In many parts of the world where there is a good solar resource, high levels of sunlight often coincide with a peak in demand for air conditioning, so solar power, particularly in the form of rooftop solar panels, can provide synchronized peak power. In addition, some solar thermal power plants can incorporate thermal energy storage, which will allow them to operate round-the-clock, depending on the size of the energy store. Otherwise, solar power is generated when the sun shines and must be fed into the grid immediately. This means that under normal circumstances, the solar power must be dispatched when it is available, while other generating plants must be ready to provide an alternative source of power when solar power is not available.

SOLAR THERMAL POWER GENERATION

Solar thermal power generation uses the sun simply as a source of heat. As discussed before, the energy reaching Earth's surface is mostly either infrared or visible radiation. A solar thermal plant can utilize the infrared and a small part of the visible spectrum. This energy is absorbed and used to raise the temperature of a heat-transfer fluid while the remainder is rejected.

The solar radiation as it reaches Earth is of relatively low intensity and it must be concentrated to be an effective source of heat. In solar thermal power plants this is carried out by the use of mirrors with the type of mirror defining the solar thermal power plant. Three types are in common use: a parabolic trough reflector, solar tower power plant, and parabolic dish solar power plant. A fourth type uses a Fresnel lens that approximates to a parabolic trough reflector.

There are two other types of solar thermal power plant. One is a solar pond, a large area of water exposed to sunlight that is designed to maintain a small temperature gradient between its upper and lower layers that can be used to drive a heat engine. This is a relatively low-technology solar thermal plant and it has been rarely used. The second, called a solar chimney, has only been tested at very small scale and never deployed commercially. The idea behind it is to exploit the heating of air by the sun to generate an updraft in a very tall chimney. Wind turbines are then placed within the chimney and their rotation generates electricity.

Parabolic Troughs

The sunlight that reaches Earth, while it can feel extremely hot, does not contain sufficient energy in the diffuse form in which it arrives to constitute the basis for a thermal power generation system. To make it useful, the sunlight from a large area must be concentrated. This can be achieved with a magnifying lens, but lenses form a relatively expensive way of concentrating sunlight. Much better is a concentrating reflector.

The parabola is the ideal shape for a solar reflector because it concentrates all the light incident on it from the sun at a single point called the focus. A complete parabola is circular; this forms the basis for a solar dish system (see later). However, there is a limit to the size of dish that can be built. For large-scale solar concentration, a trough-shaped reflector has proved more effective. If the trough is built with a parabolic cross-section, the reflector will bring the incident sunlight to focus at a line rather than at a single point—that is, a line running along the length of the trough (Figure 13.1). A heat absorption system is then

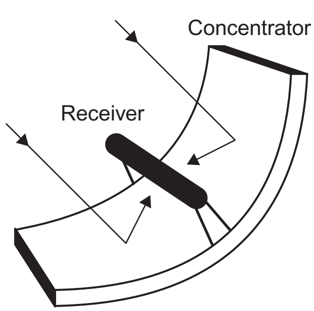

FIGURE 13.1 A solar trough system.

placed along the length of the trough at its focus. This is the basis for the solar trough, sometimes called a line-focusing solar thermal power plant.

Parabolic troughs are the oldest solar thermal technology and were first used for a plant built near Cairo in 1912 to generate steam that drove a steam engine pump. The earliest modern electricity-generating plants of this type were built in California in the 1980s and early 1990s, and nine plants are still operating. However, the economics of solar thermal technology were marginal at that time and it was not until the first decade of this century that further plants of were built.

The key element of a parabolic trough solar plant is the reflecting trough itself. These have traditionally been manufactured from 4 mm glass that is hot-formed into the required parabolic shape in a special bending plant and then mirrored. A typical trough will be up to 5–6 m wide (this is called its aperture) and will be made from individual sections up to 20 m in length that are used to build troughs up to 150 m in length. These are then mounted onto a tubular steel frame and fitted with a tracking system that enables them to track the sun across the sky. The long axis of each trough is aligned north to south so that they can rotate to track the sun from east to west. A typical 50 MW parabolic trough solar thermal plant will have 600 individual troughs and a total reflector area of around 500,000 m^2.

The cost of the reflector is the single-most important capital cost element of a parabolic trough plant, so ways of reducing costs are attracting attention. Some developers are using polymer troughs that are cheaper than glass, and lighter, but their reflection efficiency and durability are not generally as good as that of glass. Aluminum-alloy mirrors have also been used.

The heat concentrated in the parabolic trough is collected using a pipe running along the length of each mirror, at its focus, and containing a heat-transfer fluid that is pumped along the pipe. This pipe will generally be made from steel to which a special coating has been applied to maximize the heat absorption. A single heat collection pipe will run through several 150 m long mirrors to enable the heat-transfer fluid to reach a high enough temperature for power generation. These individual heat-transfer pipes are then connected in parallel in the heat collection circuit.

A parabolic reflector can concentrate sunlight between 60 and 100 times (the concentration ratio) and this is capable of raising the temperature of the heat-transfer fluid to a much as 550 °C. However, the actual temperature is often limited by the heat-transfer fluid, usually synthetic oil, which must be kept below 400° to prevent it from decomposing. For a typical plant of this type the heat-transfer fluid enters the solar field at 290 °C and leaves at 390 °C.

To generate electricity from sunlight, the heat contained within the heat-transfer fluid is used to raise steam to drive a steam turbine (Figure 13.2). This is carried out by passing it through a series of heat exchangers. Modern plants often use a reheat steam turbine that has two turbine sections with the steam heated again when it exits the first section and before entering the second section.

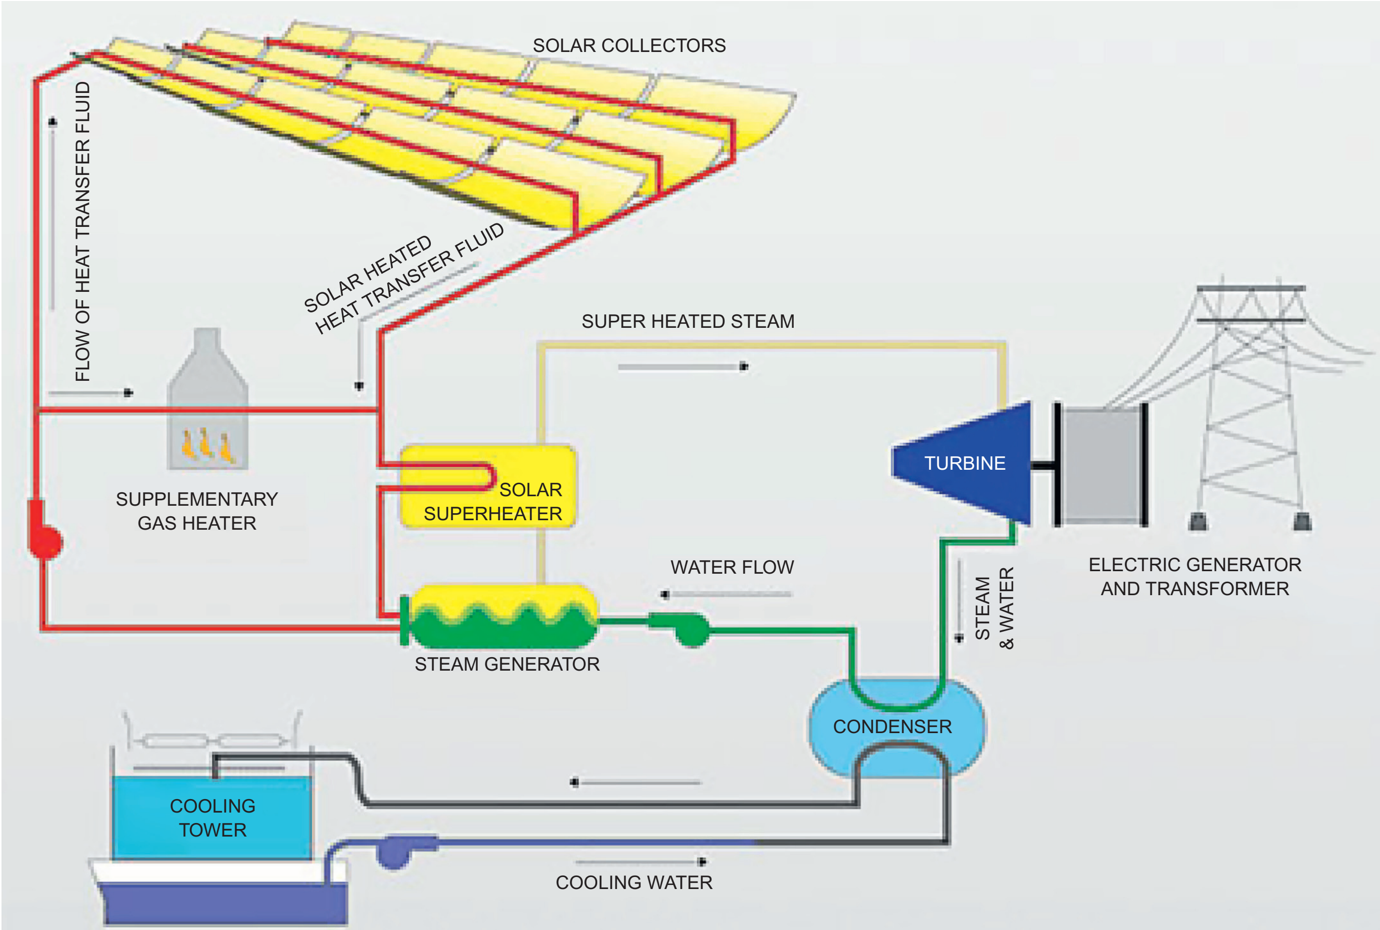

FIGURE 13.2 Schematic of a solar parabolic trough power plant.

This is usually more efficient than using a single-stage turbine. A recent 50 MW plant in Spain used a 50 MW reheat steam turbine with steam conditions of 370 °C and 100 bar.

Some plants also use supplementary heating provided by natural gas to both improve their reliability during variable sunlight conditions and to extend the range of operation into the early morning and late evening. Sunlight collection efficiency is around 75% but the low temperature of the heat-transfer fluid makes the turbine cycle relatively inefficient, at around 38%. Overall efficiency is generally 15–16%.

The operation of a parabolic trough plant can be improved by adding a heat storage facility. With this, the plant can operate for longer each day. It is not considered efficient to store the heat-transfer fluid itself, so plants that adopt this configuration usually use an alternative heat storage medium. The most common is molten salt, usually a mixture of nitrates, that can operate up to 600 °C.

A heat storage system comprises two tanks, one containing hot heat storage fluid and the other cold fluid (or relatively cold, for it will still be at around 300 °C to maintain it in the liquid state). To store energy, molten salt is taken from the colder tank and heated using the heat-transfer fluid then stored in the hot tank. To extract energy again the cycle is reversed with salt from the hot tank used to raise steam for power generation before being returned to the cold tank. With storage, the capacity factor of a solar thermal plant of this type can be raised from around 20–30% to 60%.

The amount of energy that can be stored will depend on the generating capacity of the plant and the size of the solar field. The larger the solar field relative to the size of the plant turbine, the more energy can be stored. However, since the solar field is the single-most expensive part of the plant, the economics of storage need to be weighed carefully. In principle, a plant could be designed to operate 24 hours each day, but generally they are designed to be capable of supplying power during the main periods of grid demand rather than continuously.

Since 2007 a number of commercial solar trough power plants have been built. The largest concentration of these is in Spain. Many of these installations are around 50 MW in generating capacity and a number include some form of energy storage. Larger-capacity plants have been proposed, particularly in the United States, but these have not yet been built.

In addition to the standard parabolic trough plant it is possible to build a hybrid solar fossil fuel power plant using trough technology (Figure 13.3). Such plants are based on a natural gas–fired combined cycle power plant but with an additional solar collection field. In this type of plant natural gas is burned in the gas turbine to generate power and then the exhaust gases are fed into a heat-recovery steam generator to provide steam, but with the steam production boosted during sunlight hours by heat from the solar field. These plants are normally called integrated solar combined cycle or ISCC plants. Many of the ISCC

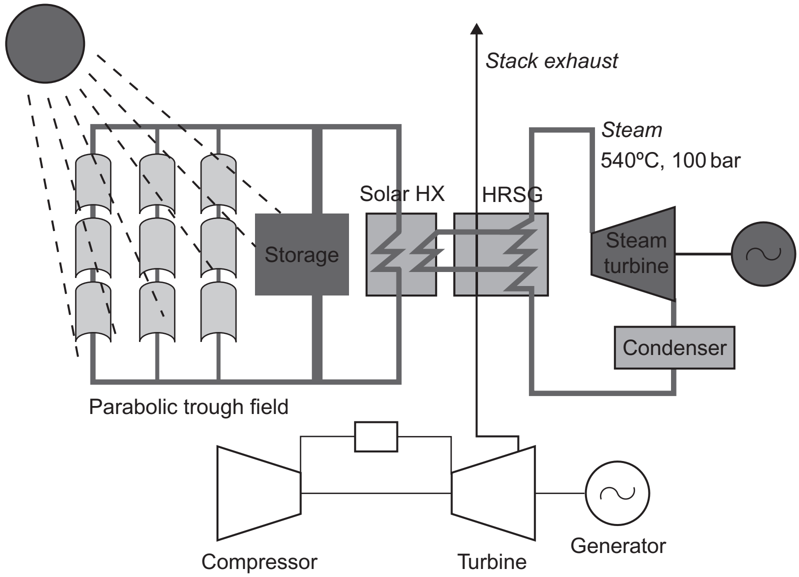

FIGURE 13.3 An integrated-solar-thermal/combined cycle power plant utilising solar troughs.

plants that have been constructed are in North Africa or the Middle East where both cheap natural gas and abundant sunlight are available, but the configuration can be exploited in many parts of the world.

Although most parabolic trough power plants use a synthetic oil as the heat-transfer fluid, the efficiency of the plants could be increased by using a direct-steam cycle. This would involve doing away with the heat-transfer fluid and heating water to generate steam directly within the parabolic trough heat collection and transfer circuit. While this presents some engineering problems, it would allow much higher steam operating conditions to be achieved, leading to higher heat engine efficiency. Direct-steam technology is being actively developed by several companies and research organizations. Another area that could add to performance would be the use of concrete for energy storage. Concrete would offer a much cheaper option than molten salt because it can store heat at a higher temperature, allowing smaller stores to be built as well as enabling a higher thermodynamic efficiency to be achieved.

Solar Towers

The solar tower takes a slightly different approach to solar thermal power generation. Whereas the parabolic trough array uses a heat collection system spread throughout the solar array, the solar tower concentrates heat collection at a single central facility. The central facility includes a large solar energy receiver and heat collector that is fitted to the top of a tower. The tower is positioned in the center of a field of special mirrors called heliostats, each of which is controlled to focus the sunlight that reaches it onto the tower-mounted solar receiver. This type of solar thermal plant is referred to as a point-focusing solar thermal plant (Figure 13.4).

The earliest solar tower was a 500 kW test rig constructed in Spain in 1981. It used liquid sodium as its heat-transfer medium (similar to some types of experimental nuclear reactor). In 1982 a project called Solar 1 was built at Barstow in California with a generating capacity of 10 MW based on a water/steam heat-transfer cycle. Solar 1 was upgraded to Solar 2 in 1996 and used molten salt to absorb and transfer the heat. Tests were carried out elsewhere too, but it was not until 2007 that the first commercial solar tower project, called Planta Solar 10 (PS10), began operating in Spain. The technology remains less well developed than solar trough technology.

The solar tower power plant is essentially an approximation of a massive parabolic dish, with a typical layout shown in Figure 13.5. The mirrors that

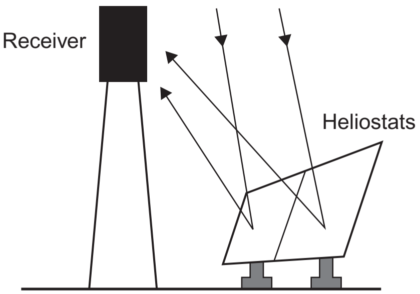

FIGURE 13.4 A solar tower system.

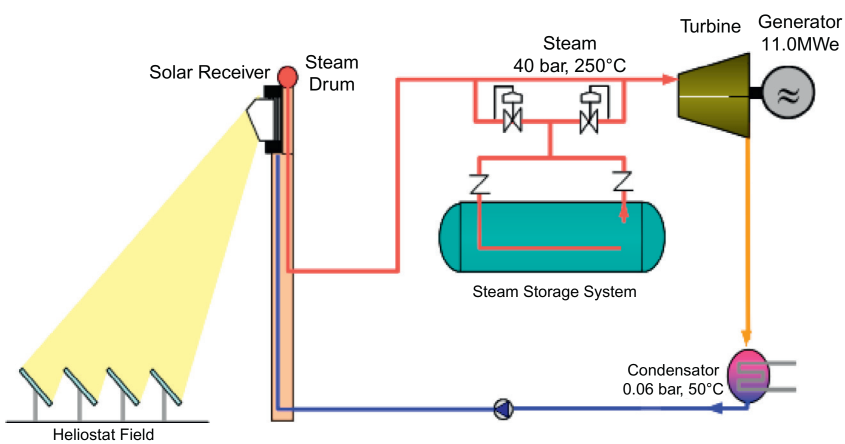

FIGURE 13.5 Schematic of a solar tower with a direct-steam energy capture system and heat storage.

make up its solar field are all parabolic reflectors that concentrate sunlight to a focus at the top of the central tower. However, each ring of reflectors belongs to a parabola of slightly different size. Each mirror must be able to track the sun independently, making the solar collection system relatively more expensive than for a solar trough plant. The advantage is that a higher temperature can be achieved in this type of plant. A concentration factor between 600 times and 1000 times is possible, allowing temperatures between 800 °C and 1000 °C to be reached. Existing plants have only achieved 600 °C, but even so this provides a 20% increase in efficiency in the steam cycle over a solar trough plant using a synthetic-oil heat-transfer fluid.

The maximum size of a solar tower power plant is determined by the size of the solar field that can be built. This will be limited because once the distance from the heliostat to the central receiver becomes too large, efficiency will drop. It appears likely that the largest single field tower will be around 200 MW or less. Larger plants will have to be based on multiple fields and towers.

The mirrors, or heliostats, used in a solar tower field are almost flat and this helps make them cheaper to manufacture than solar troughs. The nature of the solar tower means that flat terrain is not necessary either to build the plant. For a solar trough plant a flat site is essential. Heliostats are typically up to 120 m^2 in area. The 11 MW PS10 plant in Spain has 624, occupying an area of 60 ha or 5.5 ha/MW. A second Spanish plant, PS20 with 20 MW capacity, has 1255 heliostats covering an area of 90 ha.

Both or these Spanish solar towers use a direct-steam system to capture heat and generate steam. The steam conditions are relatively low at 250 °C and 40 bar and this limits the overall efficiency of the plants. Steam storage tanks are also used to extend the operating range of each plant but this is only suitable for short-term storage and operation is limited to the hours of sunlight.

An alternative approach is to use a molten salt as the heat absorption and transfer fluid. This was the approach used in the demonstration Solar 2 project in California during the 1990s. Such an approach also lends itself to heat storage since the molten salt can be stored hot in a cycle similar to that outlined earlier for heat storage with a parabolic trough plant. Hot salt is then taken from the storage tank and passed through heat exchangers to generate steam before being stored again in the cold tank. Though more expensive than a direct-steam plant, the molten salt storage system allows the plant to operate for much longer because of its higher storage capability.

There is a third design for a solar tower that could potentially offer even better performance by using air as the heat-transfer fluid. This is the basis of the volumetric air receiver, built from ceramics and capable of heating air to 1000 °C (Figure 13.6). In the simplest air system, the air is passed through heat exchangers and used to generate steam to drive a turbine. However, a more advanced cycle uses pressurized air that can be used to drive a gas turbine, once heated. The hot air exiting the gas turbine is then used to generate steam for a steam turbine in a cycle similar to the combined cycle power plant. An air cycle

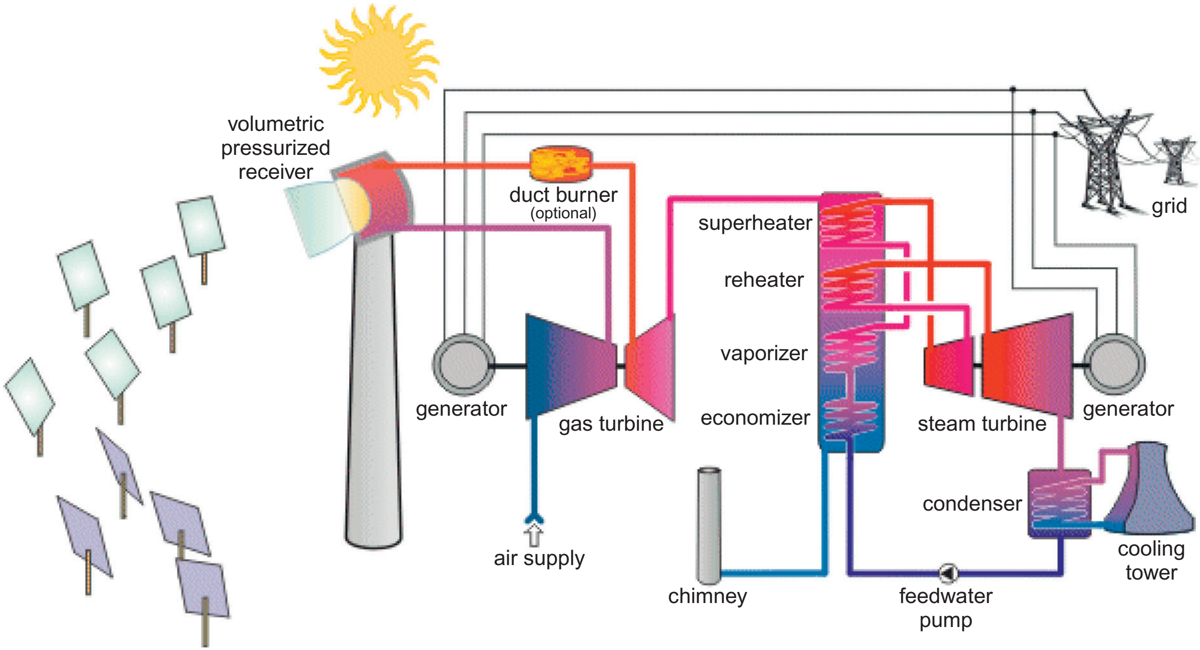

FIGURE 13.6 Schematic of a pressurized volumetric air solar tower receiver system.

is not easily adapted for energy storage but can be fitted with supplementary natural gas firing if necessary.

Spain is the location for the largest solar towers in operation today. These include the PS20 20 MW plant and the Gemasolar plant with a generating capacity of 19.9 MW. The latter, which began operating in 2011, uses a molten-salt heat-transfer fluid and is sized so that it can provide power for up to 15 hours each day without the need for sunlight. This allows it to claim 24-hour operation. Overall plant efficiency is estimated to be 15% and the plant has a capacity factor of 65%. A much larger project is under construction at Ivanpah in California. This plant will comprise three facilities with a combined generating capacity of 377 MW. Each facility will have several towers, with the individual tower size expected to be 33—50 MW. These are based on a direct-steam cycle with steam conditions of 550 °C and 160 bar. The first unit started delivering power to the grid in 2013 from the first of three towers. These plants require 3.4 ha/MW.

Solar Dishes

A solar dish power plant is typically a single parabolic reflector similar to a satellite antenna but with a heat engine instead of a microwave receiver at its focus (Figure 13.7). The dish and its engine are constructed as a single unit, mounted on a framework that will allow it to track the sun across the sky. Most solar dish power units are 25 kW or less though bigger ones have been built. Deployed singly, they can provide a self-contained power supply unit. However, they can also be deployed in larger numbers to provide higher generating capacity.

Most solar dishes use simple self-contained heat engines to generate power, but in one or two recent projects solar dish systems have been designed to provide steam for a steam turbine. The main advantages of solar dishes are their simplicity and efficiency. Single solar dishes are capable of achieving 30% energy conversion efficiency, much higher than any of the other solar thermal technologies.

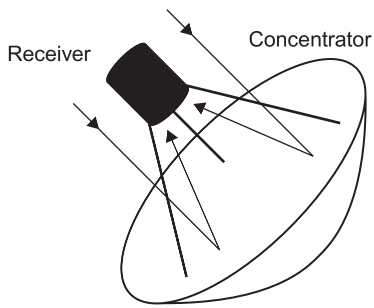

FIGURE 13.7 A solar dish.

The first solar dish energy devices were built by the French mathematician August Mouchet in the 19th century to drive a steam engine. The modern version, designed for power generation, began to appear after the 1970's energy crises. However, the technology has struggled to prove itself economically and it was only at the end of the first decade of the 21st century that a demonstration project involving a large array of solar dishes finally started generating power.

The solar dish can be considered to be a solar tower in miniature, a point focusing solar thermal system with a high concentration ratio. With its single, continuous reflector a solar dish can achieve a concentration ratio of 2000. This allows a very high temperature up to 1000 °C to be achieved, allowing high energy conversion efficiency. The heat engine at the center of a dish is matched to the heat input, with a 10 m diameter dish capable of supplying around 25 kW of electrical power with a direct solar insolation of 1000 W/m^2. Units capable of providing up to 100 kW have been built, but at larger sizes the dishes become massive and difficult to manage.

As with other types of solar thermal power plants the reflectors for solar dishes are generally made from mirrored glass (Figure 13.8). However, a single 10 m diameter glass reflector would be prohibitively expensive (and heavy) so they are usually made up from an array of small elements that approximate the large dish. Plastic and aluminum reflectors have also been tested, as well as a novel technique in which a thin plastic membrane is pulled taut over a hoop and then deformed by applying a partial vacuum to one side of it.

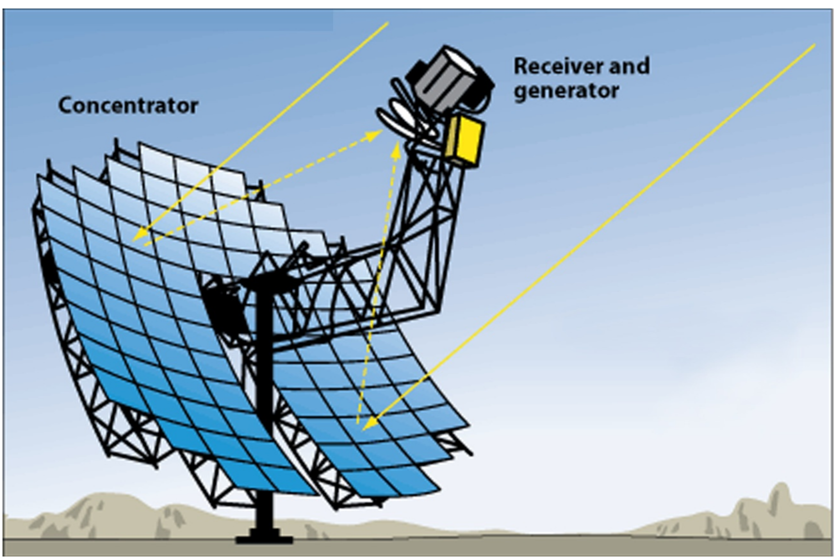

FIGURE 13.8 A solar dish power system with multiple mirrors.

To achieve high efficiency, a solar dish must track the sun accurately. For larger dishes this is usually achieved by tracking on two perpendicular axes. For smaller dishes polar tracking is often deployed. This involves rotating the dish about an axis parallel to Earth's axis of rotation at 15°/hour to match the speed Earth's rotation. Movement about a second axis is often used to take account of the seasonal movement of the sun in the sky.

The generation of electricity with a solar dish is normally carried out with a small self-contained engine, usually either a Stirling engine or a micro-turbine fixed at its focus. In both cases the heat at the focus of the dish is used to heat a thermodynamic fluid that drives the engine, generating power. In a Stirling engine that fluid is hydrogen or helium held within the engine in a closed cycle. With a micro-turbine the thermodynamic fluid is air and the cycle is that of a gas turbine cycle. The Stirling engine is capable of very high efficiency and, although operating engines in solar dishes have only achieved just over 31% efficiency, 40% may be feasible. The micro-turbine is less efficient, with a potential efficiency around 30%, but they are cheaper. In addition, the heat input into the micro-turbine can be boosted with natural gas if necessary. That is less easy to achieve with the Stirling engine, which has a completely closed cycle.

It is possible to build a solar dish that can use the heat it collects to generate steam for a steam turbine. To make this effective it will normally be necessary either to construct one very large dish or build a series of smaller dishes connected to a single heat circuit. Dish systems based on a steam cycle have been developed by the Australian National University.

The commercialization of solar dish technology has been slow. The first semi-commercial project, the Maricopa power plant in Phoenix, Arizona, entered service in 2010. The plant has 60 parabolic dishes with an aggregate generating capacity of 1.5 MW. Its construction was supported by the U.S. government. A project of similar capacity at the Tooele Army Base in Utah began testing in the middle of 2013. Much larger schemes have been proposed in California.

Fresnel Reflectors

The Fresnel reflector represents an attempt to reduce the cost of solar thermal technology by creating a cheap heat collection system. In this case the model is the parabolic trough. However, instead of a single parabolic trough, the Fresnel system employs a series of long, almost flat reflectors that approximate the trough. Each individual reflector is mounted so that it can track the sun and the energy it collects is focused onto a stationary pipe above the reflectors running the length of the trough. Figure 13.9 provides a typical layout.

The earliest solar thermal plant based on Fresnel reflectors was developed by an Italian, Giovanni Francia, who patented the system in 1962 and built a prototype in France in 1964. Further development took place in Australia in

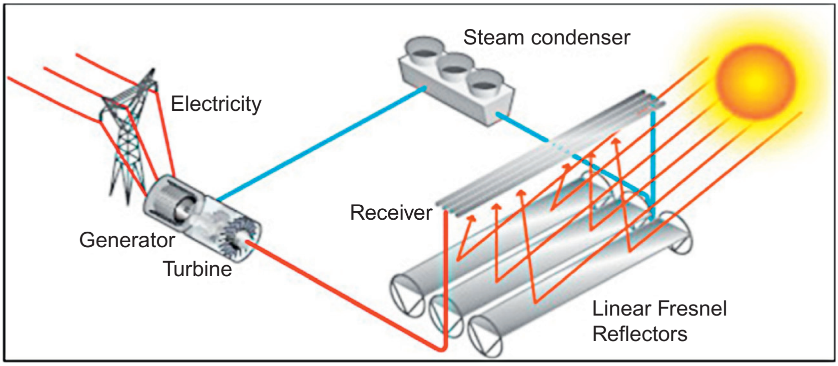

FIGURE 13.9 A Fresnel lens thermal power plant.

the 1990s where a large demonstration scheme was constructed, and since then plants have been built in the United States and in southern Spain.

The Fresnel system bears some similarity to the solar tower approach with a series of flat reflectors approximating the parabolic reflector, although in this case the reflector it mimics is a trough. Like the tower, these reflectors can be mounted close to the ground, simplifying construction and increasing stability while the reflector and heat collector are separated, making the design simpler. The reflectors are long and either flat or slightly curved, with around 20 replacing a single trough. As with a trough, each reflector is made of around six individual segments. They are usually made from glass but other materials are possible. Collection efficiency is lower than for a parabolic trough, with a concentration ratio of between 30 and 50 typical.

Fresnel power plants usually use direct-steam generation to improve efficiency and reduce costs. They are claimed to have smaller space requirements than traditional trough plants with around 1.3 ha/MW. One of the largest operating plants is the Augustin Fresnel 1 plant in France with a generating capacity of 250 MW. This plant uses water as the heat-transfer fluid, the latter exiting the heat collection field at 300 °C. The turbine operating pressure is 100 bar. The plant began generating power for the grid in 2012.

Other Solar Thermal Technologies

The solar chimney (sometimes confusing since also called a solar tower) uses heat from the sun to raise the temperature of an enclosed volume of air, creating a pressure that can be used to produce an updraft in a tall chimney. Turbines within the chimney can then generate electricity.

To make this feasible, a large circular greenhouse must be constructed to capture heat and raise the temperature of the air enclosed within it, as shown in Figure 13.10. The outer edges of the greenhouse are open so that air can be drawn in. Meanwhile, at the center of the circle rises a tall chimney. The hot air within the greenhouse is drawn toward and up this chimney, creating an air flow that can be exploited by wind turbines to produce electricity.

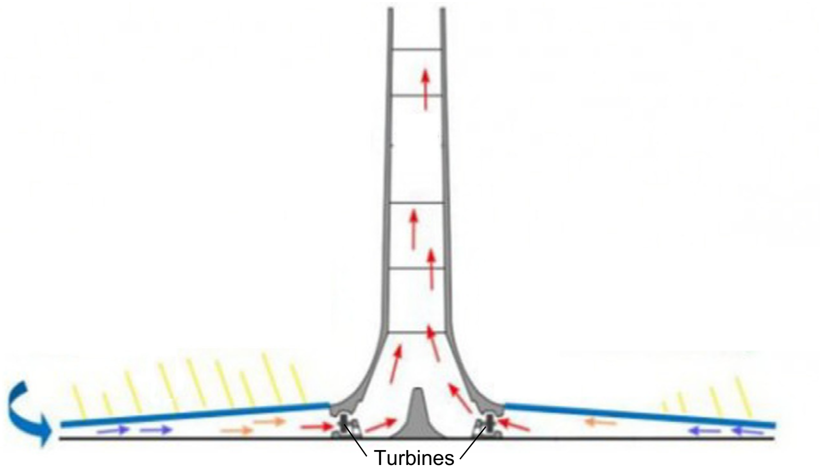

FIGURE 13.10 Schematic of a solar chimney.

Although the idea for a solar tower was published as early as 1903 by a Spaniard, Isidoro Cabanyes, the first to be built was constructed in 1981 in southern Spain by a German company. This had a 180 m high, 12 m diameter tower and housed a single 50 kW turbine to test to principle. The project operated until 1989. The concept has since been adopted by a U.S. company that has put forward plans to build a major plant of 100–200 MW but no plant has yet been constructed. In 2010 a 200 kW tower started operating in Inner Mongolia. There are plans to expand this Chinese development to around 27 MW.

A solar pond is also a relatively old idea that has rarely been exploited. It consists of a shallow pond, usually between one and two meters in depth, filled with brine and covered with some form of membrane to prevent evaporation (Figure 13.11). When sunlight heats the brine it creates a layer of hot, concentrated brine at the bottom of the pond. As a consequence of its density the brine remains at the bottom of the pond even though it is hotter than the water above it.

The hot brine can reach 70–80 °C while the upper layer will be between 20 °C and 30 °C, creating a temperature gradient that can be used to drive a small heat engine. This is usually a small organic Rankine cycle turbine that can extract energy from such a tiny temperature difference. The technology was tested in Israel during the 1980s when a 5 MW project was built in the Dead Sea region, but it does not appear to have been pursued, either there or elsewhere.

PHOTOVOLTAIC DEVICES

Photovoltaic devices, often called solar cells, are solid-state devices that are capable of absorbing sunlight and converting it into an electrical current. The earliest observation of the photovoltaic effect was by the French scientist Antoine-César Becquerel but the practical solar cell was developed by scientists at Bell Laboratories in the United States in the 1940s and 1950s. These cells

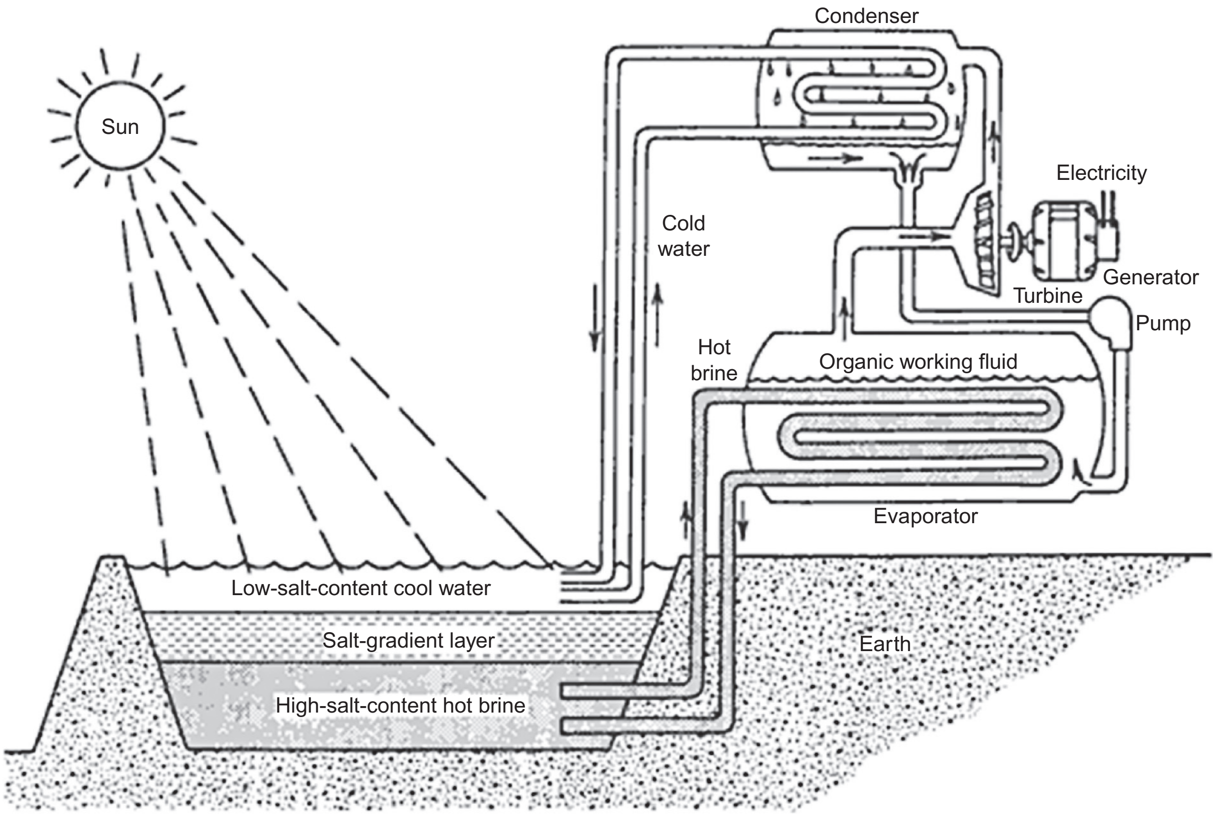

FIGURE 13.11 Schematic of a solar pond power plant.

were around 6% efficient and during the 1960s they were used in the U.S. space program. As costs came down they began to be used in the 1980s for remote terrestrial applications and during the 1990s, with prices continuing to fall, grid-connected solar cells began to appear. Their use expanded dramatically from the end of the 1990s, and by the beginning of the second decade of the 21st century they had become the third most important source of renewable electricity after hydropower and wind power. Their simplicity, combined with continued lowering costs, means that they could overtake wind power in the next decade or two.

Solar Photovoltaic Technology

The operation of a solar cell depends on a fundamental property of a semiconductor called its bandwidth. The bandwidth is a function of the semiconductor's atomic structure that results in an energy gap between the top layer of full electron energy levels and the first set of empty energy levels. The energy gap is too large for electrons to jump from the lower to the upper as a result of thermal activation. However, electrons in the lower level can become promoted from the lower to the upper level, across the bandwidth, by absorbing photons of electromagnetic radiation. For this to be possible, the photon must contain at least as much energy as the size of the energy gap between the two sets of energy levels in the semiconductor.

The range of electromagnetic radiation that the cell can absorb is determined by the size of its bandwidth. Semiconductors that are useful for solar cells have bandwidths that make them capable of absorbing photons within the visible region of the solar spectrum, but any photon with an energy lower than the bandwidth (e.g., infrared radiation) will not be absorbed. However, all those with energy greater than the bandwidth can be absorbed. Table 13.2 shows the bandwidths of some semiconductors that are commonly used for solar cells.

Each photon of light energy that is absorbed by the semiconductor is captured by an electron within the material. In absorbing the energy, the electron acquires an electrical potential (it has more energy than most of those around it). This potential can be made available as electrical energy, as an electric current.

TABLE 13.2 Bandwidth of Some Common Solar Cell Semiconductors

Semiconductor	Bandwidth (eV)
Silicon	1.11
Cadmium telluride	1.44
Gallium arsenide	1.43
Copper indium gallium diselenide	0.9–1.7

The current is produced at a specific fixed voltage called the cell voltage. The cell voltage is, again, a property of the semi-conducting material. For silicon it is around 0.6 V.

The energy contained in light increases as the frequency increases from infrared through red to blue and ultraviolet light. However, a solar cell must throw away some of these frequencies since it can only absorb light above its cell threshold. Light that is of an energy below this threshold simply passes through the material.

It might seem sensible, therefore, to use a semiconductor with a low threshold or bandwidth. However, this will lead to a cell with a low output voltage since this is also directly related to the threshold for absorption. There is another drawback to using a semiconductor with a small bandwidth. When a photon of light with energy much greater than the threshold energy is absorbed, it loses all the energy above the threshold value. The surplus energy is essentially thrown away (it emerges as heat) and cannot be used for electricity production.

These two factors mean that the lower the bandwidth and absorption threshold, the more energy is thrown away as heat, but the higher the bandwidth, the more energy simply passes through the material without being absorbed. The optimum is, therefore, a compromise between the two competing effects.

The optimum bandwidth for a solar cell semiconductor is 1.43 eV. The bandwidth of silicon, at 1.11 eV, is below the optimum but it has proved the most effective solar cell material to date and has the largest market share. Silicon is used in three different forms: crystalline, semi-crystalline, and amorphous. Crystalline silicon is the most efficient but also the most expensive to produce, while amorphous silicon is both the least efficient and the cheapest. Alternatives to silicon include cadmium telluride, which is cheaper to produce and is always in amorphous form. Its bandwidth is very close to the optimum. Other materials such as gallium arsenide are also more efficient that silicon, particularly in crystalline form, but are much more expensive.

Types of Solar Cell

Solar cells are manufactured using technologies similar to those used to manufacture microchips and transistors. Most of these are made using slices of perfect silicon crystals that are then etched and doped to create the complex structures that are required for computers and other electronic devices. Solar cells, though normally simpler in structure than a microchip, can be manufactured in a similar way.

Silicon solar cells made from single crystal silicon are the most efficient available with reliable commercial cell efficiencies up to 20% and laboratory efficiencies measured at 24%. Even though this is the most expensive form of silicon it remains the most popular by a wide margin due to its high efficiency and durability. Polycrystalline silicon is cheaper to manufacture but the penalty is lower efficiency with the best measured around 18%. Cheapest of all to produce is amorphous silicon, which can be made using vapor deposition techniques rather than by expensive crystal growing. However, its best efficiency is only 8%. Amorphous silicon also suffers from degradation when first exposed to light, a feature not seen with crystalline material. This can reduce its initial efficiency by up to 20%.

All silicon solar cells require extremely pure silicon. The manufacture of pure silicon is both expensive and energy intensive. The traditional method of production required 90 kWh of electricity for each kilogram of silicon. Newer methods have been able to reduce this to 15 kWh per kilogram. This still means that depending on its efficiency, a silicon solar cell can take up to two years to generate the energy needed to make it. This compares with around five months for a solar thermal power plant. Manufacturers of crystalline silicon are concentrating on ways of reducing the cost of crystalline material by cutting it more efficiently or by finding new ways of growing it. This is helping push costs down, and crystalline silicon remains competitive in spite of efforts by thin film manufacturers using much cheaper materials.

Another crystalline material used for solar cells is gallium arsenide. This has an almost perfect bandwidth for a solar cell and in the laboratory efficiencies of 28% have been recorded. However, practical cells only reach 20% and the material is both expensive and composed of hazardous materials. It is rarely used, except for special applications.

The main alternative to crystalline silicon for solar cells is some form of thin film. From a manufacturing point of view these are attractive because they can be produced using cheap techniques such as vapor deposition or even printing. Amorphous silicon is one alternative but it is not as cheap to produce as cadmium telluride and the latter has a much higher efficiency, with the best recorded at 17% though commercial cells rarely achieve more than 11%. This material also has an almost optimum bandwidth for a solar cell (1.44 eV) and its potential efficiency could approach 30%. Cadmium telluride is also attractive because it is possible to produce solar cells on a variety of substrates including building components and flexible plastic sheets.

The maximum efficiency possible with a single-layer solar cell of any semiconductor is 33.7%. It is possible to build more efficient solar cells by layering cells one on top of the other. Such multilayer or multijunction cells place the semiconductor with the largest bandwidth at the top. This will absorb high-energy radiation but that of lower energy will pass through it to reach the layers below where further semiconductor layers of lower bandwidth are placed. In principle, it is possible to create a cell with up to 50% energy efficiency with multilayer cells, but they are much more expensive to manufacture. The best recorded efficiency achieved to date is 43.7%.

Cell Structures

The conventional structure for a solar cell is planar (Figure 13.12). The cell is made of a thin layer of semiconductor, the top surface of which is doped with an

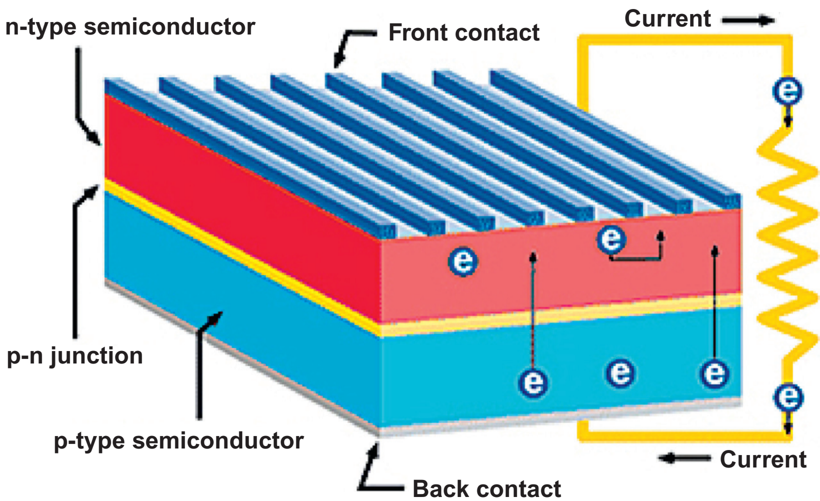

FIGURE 13.12 Solar photovoltaic principle.

impurity to create a *p-n* junction within the body of the material. It is this junction, with a built-in voltage gradient, that captures electrons once they are generated by light absorption and sweeps them away to form an electric current. To collect the current, electrical contacts must be formed onto the semiconductor. The rear of the cell can be covered with a planar collector, but on the front surface, where light is absorbed, the collector area must be minimized or it will interfere with light absorption. The front collector is usually made from narrow fingers of metal that allow the maximum amount of sunlight to strike the semiconductor surface. The semiconductor will often be placed on a substrate to give it additional strength and the top layer will be covered with a protective coating.

More advanced structures are possible. The front contacts can be buried, edgeways, into the surface of the cell to minimize the shadow effect of the contact that reduces efficiency. More modern cells have also been developed that move both contacts to the back of the cell by allowing the front doping to be carried to the back of the semiconductor. This structure is more complex but creates a more efficient cell.

Efficiency depends on light absorption. This can be improved by making the front surface of the cell nonreflecting and by making the back reflective so that any light reaching the back layer is reflected back into the semiconductor. Such developments mean that the number of stages in the fabrication of a solar cell can increase to perhaps 15 from a more usual 6–8. However, the lowering of production costs and the increased efficiency are making such developments worthwhile.

Such advanced techniques are generally only used with crystalline silicon. Amorphous or thin-film devices rely on cheaper fabrication techniques for their economics. Cadmium telluride, the most popular thin-film material, is usually formed on a substrate such as glass. However, it can also be deposited onto flexible plastics. These thin films can be produced in much larger areas than crystalline silicon, which helps make them more cost effective. Copper-indium-gallium-diselenide is another thin-film solar cell material. This has the advantage that by varying its composition, the bandwidth can be varied, which offers the potential to tailor cells to particular applications or to manufacture complex multilayer cells with specific characteristics.

Concentrating Solar Cells

Concentrating solar cells adopt a different strategy to planar cells. Whereas the latter capture sunlight over a large area of semiconductor, concentrating cells use cheap optical systems to concentrate the sunlight first, in a manner similar to most solar thermal technologies, before converting it into electricity. In principle, an optical system can be manufactured more cheaply than a large area planar solar cell. The device can then use a very expensive, but highly efficient,multijunction solar cell at the focus to convert the light into electricity. As noted before, this type of solar converter has achieved the highest efficiency of any solar cell at 43.7%.

The disadvantage of the concentrating cell is that it requires direct sunlight and cannot operate effectively with diffuse sunlight. While these converters have a limited market share today, there are signs that they may be attractive for utility-scale applications in high insolation regions where their high efficiency is likely to be an advantage.

Third-generation Solar Cells

In the market for solar cells, crystalline silicon cells are often referred to as first-generation solar cells and thin-film devices as second-generation cells. Third-generation cells are new organic and dye-sensitized solar cells.

Organic semiconductors operate in exactly the same way as conventional ones, but with the disadvantage that the organic material does not conduct electricity as readily as a traditional semiconducting material. The key to developing organic materials is to find an efficient way of extracting the electrical power once it has been generated. (If the excited electrons are not swept away quickly, they will simply fall back to the lower energy level in the semiconductor and lose their energy.) While this can be a drawback, the attraction of organic materials is the simplicity of their production. It is expected to be possible to print the semiconductor onto substrates, making it extremely flexible in the way it can be used. However, efficiencies of organic semiconductors are currently very low with the best laboratory value of 6% and the best production efficiency of 4%.

Dye-sensitized solar cells use special dyes to absorb light and then transfer the electron produced during the absorption to a substrate semiconducting material. These devices have been compared to photosynthesis, but the complex series of exchanges required to make them work also makes them appear similar to electrochemical devices such as batteries. As with organic semiconductors these are attractive because of the simplicity of manufacture as they too can be printed. Theoretical maximum efficiency is 31% and practical efficiencies of 13% have been achieved. Both types can maintain their outputs under low light conditions but long-term stability is still questionable.

Modules, Inverters, and Panels

Most solar photovoltaic energy conversion is based on planar solar cells but the actual solar cells account for only a part of the complete system required to create a solar energy conversion installation. Individual silicon solar cells only produce one of two watts of electricity each at, for a silicon cell, 0.6 V. Therefore, cells must be connected in series and in parallel to create modules capable of both higher power and of generating at a higher voltage (Figure 13.13). These modules must then be encapsulated to protect them from the environment. An encapsulated module of this type is called a solar panel and it may have an output of several watts to several hundred watts. Current encapsulation techniques are capable of providing modules with lifetimes of 25–30 years. Advanced encapsulation techniques may be able to extend this to perhaps 50 years of more in the near future.

The output of the solar panel is direct current. This can be used without conditioning in some remote applications, but for grid connection the DC must be converted into alternating current at the grid voltage. This is carried out using a solid-state inverter. Modern inverters, particularly for larger solar panel installations, are sophisticated devices that can provide two-way communication with transmission and distribution system operators. Other parts of the complete

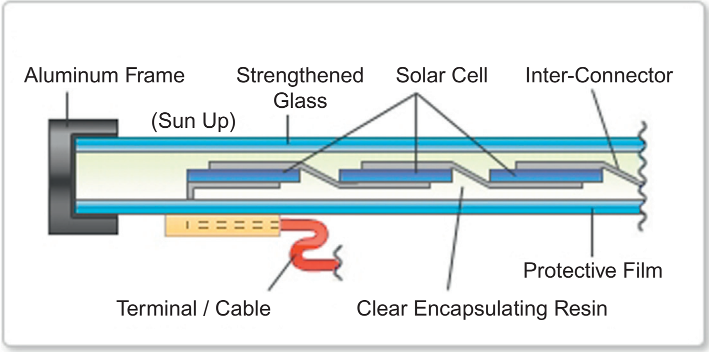

FIGURE 13.13 Cross-section of a solar cell module.

installation include the supports, typically of steel. For a modern solar photovoltaic installation the solar cell probably accounts for 50–60% of the total cost, but this is falling as fabrication costs fall.

System Types

There are broadly four types of solar photovoltaic installation, each serving a different market sector. The most important of these today is the residential sector that includes installations up to 20 kW on domestic rooftops (Figure 13.14). This is the sector that has been the focus of incentive schemes in countries such as Japan, Germany, the United Kingdom, Italy, and the United States, and it accounts for by far the largest installed capacity of solar cells, worldwide.

Allied to this is the commercial sector. This covers larger building installations at schools, hospitals, office buildings, shopping centers, and other similar organizational buildings. Installations for these types of buildings may be up to 1 MW in generating capacity.

The utility sector is the third major market sector. This involves large arrays feeding power directly into the grid. Plants of this type have capacities between

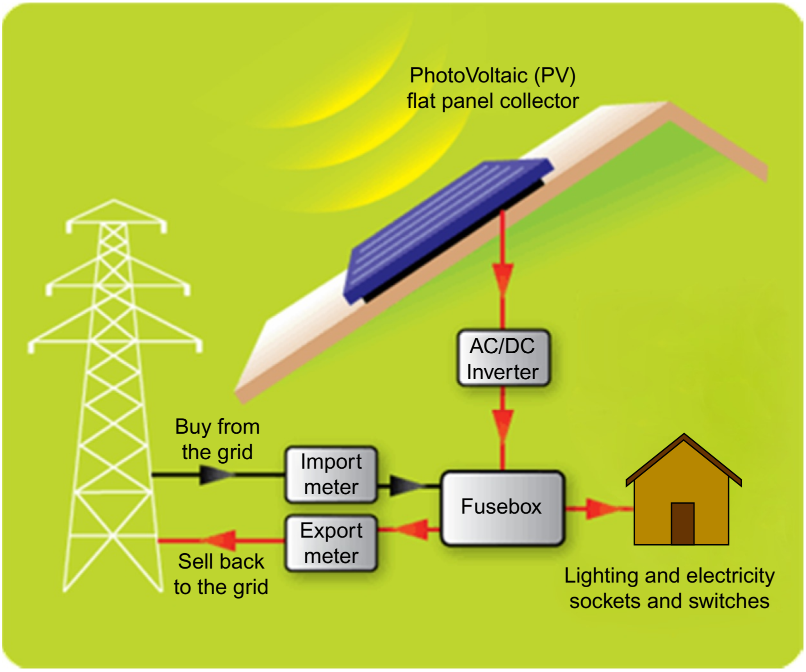

FIGURE 13.14 A grid-connected rooftop solar panel.

1 MW and 200 MW and larger installations are likely in the future. This is developing into one of the main markets for concentrating solar cells too.

The final market sector is for off-grid installations. This can range from small domestic installations to much larger utility-sized installations providing power for remote communities. In 2010 the residential sector accounted for 63% of the total market for solar cells and the utility sector for 19%. Commercial systems took 11% and off-grid systems the remaining 7%. The residential sector is expected to continue to dominate the market but the utility sector will grow strongly too over the next decade.

Solar Photovoltaic Generation and Energy Storage

As with solar thermal generation, there is a significant attraction in being able to combine solar cells and some form of energy storage. This can both extend the period during which the solar system can supply power and, where grid connected, provide a more reliable and therefore more valuable electricity source.

The most popular way of adding storage to a solar cell system is by means of batteries. Since the output of a solar cell is direct current and a battery is a DC device, the solar cell can be used to charge a battery system without any conditioning. In the past lead-acid batteries have been popular as part of off-gird solar-generating systems. Today there is a move toward using battery storage with grid-connected systems too so that the owners can take more power from their installation; several manufacturers are developing lithium-ion batteries to serve this market. Using a battery system with a grid-connected solar system is particularly attractive if the price obtained for power exported to the grid is much lower than the cost of buying power from the grid.

Battery storage is probably the easiest to implement with a solar generator but other approaches are possible. Small-scale heat storage, though relatively inefficient, can be used to harness excess solar power. For larger-scale production, systems such as hydrogen generation are feasible and even methane production.

As with all storage, much will depend on the relationship between the owner of the solar generator and the grid. It will always be more efficient, overall, to export surplus power to the grid where it can be used in a number of different ways than it will be to store it locally. Grid-level storage will generally be more efficient than local storage too although installation costs are high and this has tended to dissuade utilities from building grid storage in the past.

COST OF SOLAR POWER

Solar-generated electricity has traditionally been considered one of the most expensive sources of electricity, particularly from solar photovoltaic devices. The growth in solar thermal–generating capacity is beginning to bring costs of these technologies down, but the biggest change in the past five years has been a dramatic fall in the cost of solar cells. This has been driven by technological advances and economies of scale as production volumes have risen, but one of the main factors has been fierce competition from manufacturers in countries, such as China, undercutting the prices from traditional manufacturers.

Solar Thermal Costs

The capital cost of operating solar thermal power plants varies widely and it is difficult to achieve a consensus on the realistic cost of this type of project. This is partly a consequence of the fact that many of them are either demonstration projects or early commercial projects, so costs will be both higher than for fully commercial rollout and extremely variable.

Solar trough power plants are the most widely used solar thermal plants and their costs are the most stable. Typical costs for existing plants are around €6000/kW in Spain and £4000/kW in the United States. Solar towers based on direct-steam systems have a cost of $4000/kW to $6000/kW, but the Gemsolar molten-salt system with extensive heat storage costs around $10,000/kW. The demonstration Maricopa solar dish scheme cost $13,000/kW, although single-dish projects appear to be much cheaper. Meanwhile, one Fresnel scheme in the Unied States has a claimed capital cost of $3000/kW.

Table 13.3 contains estimated costs of various solar thermal technologies from organizations such as the U.S. Energy Information Administration, Lazard, and the U.S. Department of Energy. These suggest the cost of parabolic trough plants to be in the region of $4500–5800/kW. Solar towers are slightly more expensive at $4800–6300/kW. The solar dish capital costs in Table 13.3 are those already quoted earlier: $3300–10,000/kW.

Estimates for the levelized cost of electricity from these plants vary widely. The cheapest is between $130/MW and $150/MWh based on estimates from Lazard. Other sources put the levelized cost of power at $260–270/MWh.

TABLE 13.3 Cost Estimates for Solar Thermal Power Generation

Solar Thermal Technology	Capital Cost ($/kW)	Levelized Cost of Electricity ($/MWh)
Parabolic trough	4500–5800	150–272
Solar tower	4800–6300	130–260
Solar dish	3300–10,000	n/a

Source: U.S. Energy Information Administration, Lazard, and U.S. Department of Energy.

Solar Photovoltaic Costs

Solar cells were traditionally considered to be one of the most expensive sources of electricity, but the fall in costs in the past five years means that there is now talk of them being able to compete with other technologies without subsidy by the end of the second decade of the 21st century.

The cost of a crystalline silicon solar cell in Europe in 2009 was €2.62/W according to pvXchange, but by the beginning of 2013 the cost had fallen to €0.79/W. This is a result of both greater competition and larger volumes. Greenpeace and the European Photovoltaic Industry Association have estimated that solar cell costs have fallen by 22% each time the installed capacity has doubled. Similar trends can be seen for other solar cell types such as cadmium-telluride thin films.

While spot prices fall, module costs have fallen too but they vary substantially from region to region. In Germany, for example, the typical module price for small rooftop systems was €1.8/W ($1800/kW) in 2012 compared to perhaps €0.5/W in China. This is still only part of the cost of a rooftop installation but it is probable that the installed cost of large solar photovoltaic systems has fallen below €2000/kW in some countries.

The rate at which installation costs are falling is outstripping most predictions from energy organizations. There is scope for prices to fall much further and they could be below €1000/kW by 2025. The rate of change in prices also means that most estimates of the levelized cost of power from solar cell installations are much higher than actual costs. For example, the 2012 edition of the U.S. Energy Information Administration's Annual Energy Outlook put the cost of electricity for a solar photovoltaic power plant entering service in 2018 at $144/MWh.[1]

1. Levelized Cost of New Generation Resources in the Annual Energy Outlook 2013, U.S. Energy Information Administration.

Chapter 14

Marine Power Generation Technologies

In the search for sources of electric power to replace fossil fuels, technologists have explored all the natural phenomena that might be exploited to generate electricity. As a consequence, wind power, solar power, and biomass-based generation are today becoming mature technologies alongside hydropower. However, one potential source of energy has yet to exploited commercially—the energy contained in the seas and oceans.

Part of the reason for the slow advance of marine power generation is the hostile environment in which marine energy conversion devices have to operate. The development of power generation technologies that can withstand these conditions requires great tenacity and it has not always been clear that the seas can yield an economical means of generating power. In spite of this technologists have persevered, and while marine power generation remains at the development and demonstration stage, there is greater confidence today than there has ever been that it will produce results in the near future.

There are four ways in which energy can be extracted from the seas to provide power generation. Perhaps the most significant today, because it is the easiest to exploit, is the energy provided by marine currents. These currents are generated by tidal movements and are normally found in estuaries and around rock formations close to shore. Energy can be extracted from them by using underwater turbines that are conceptually similar to wind turbines.

Sea and ocean waves offer the second source of energy. Waves are generated by the wind and like the wind they are erratic, but it is possible to exploit the energy they contain with devices of various kinds, some mounted close to the shore, some in deep water, some floating, and some anchored to the seabed. This is potentially the most difficult marine energy source to tap but there is a significant resource available if it can be mastered.

The third source of marine energy is provided by the sun. The world's seas, particularly those in the tropics, are massive solar collectors absorbing heat energy throughout the day. This absorption creates a layer of hot water at the sea's surface, while below the surface the water temperature is much lower, creating a temperature gradient that can be exploited to drive a heat engine. The technique is known as ocean thermal energy technology (OTEC).

The fourth source of energy is probably the most unexpected. When fresh water from a river reaches the sea and mixes with salt water, energy is released. This energy of dilution is predicted by classic thermodynamics, and in the past decade one or two pilot projects have sought, with some success, to take advantage of this as a way of producing electricity. Today this resource is referred to as salinity gradient, and while it is immature, it does offer future promise.

MARINE ENERGY RESOURCE

To arrive at an estimate for the size of the global marine energy resource it is necessary to make assumptions about the actual amount of energy that each type of resource contains and the efficiency with which it can be extracted. Producing such estimates can be a challenge for resources as widely distributed and as difficult to define as these, but attempts have been made. Some are shown in Table 14.1. The table contains figures for the annual energy that might be extracted from each resource based on figures from the International Energy Agency (IEA), and figures for the potential generating capacity each might yield, based on research from Powertech Labs.

The largest resource based on the IEA figures is wave power that could, in principle, provide up to 80,000 TWh of energy each year. OTEC is the next largest with a potential 10,000 TWh/year, followed by salinity gradient with 2000 TWh/year, and finally tidal current that could yield 800 TWh/year. Other estimates provide significantly different values so these should be used for broad guidance only. In terms of the generating capacity each might support, wave power is again the largest, potentially, with up to 10,000 GW. Tidal current might provide 5000 GW, OTEC 2600 GW, and salinity gradient 2000 GW. Based on these evaluations, even the smallest could potentially provide a significant contribution to global power generation.

Of the four resources, tidal stream has the smallest potential, but because the currents that make up the resource are located close to shore, it is perhaps the

TABLE 14.1 Marine Energy Resources

Marine Energy Resource	Potential Annual Energy Available	Potential Generating Capacity
Tidal current	800 TWh/y	5000 GW
Wave power	80,000 TWh/y	1000–10,000 GW
Ocean thermal energy technology	10,000 TWh/y	2600 GW
Salinity gradient	2000 TWh/y	2000 GW

Sources: *Ocean Energy, Opportunity, Present Status and Challenges*, International Energy Agency/Ocean Energy Systems, 2004–2006; and G. Bhuyan, *Ocean Energy: Global Technology Status, Opportunities and Challenges for Canada*, Powertech Labs Inc., 2005.

easiest to exploit and so the generating capacity it might provide is relatively high. The existence of tidal currents depends critically on local topologies and some regions have much better regimes than others. The U.K. coastline is estimated to hold around 40 TWh/year, and in Alaska perhaps 100 TWh/year could be exploited.

There are other ocean currents that are not caused by the tides, such as the Gulf Stream that operates in the Atlantic Ocean. This could perhaps be exploited too using the technology for tidal stream energy capture. Estimates from Florida State University suggest that 8 GW of generating capacity could be installed to extract energy from the Gulf Stream as it passes the Florida coast. Capturing energy from the Gulf Stream in the deep ocean would be more difficult.

Waves are generated when the seas absorb energy from the wind. The longer the reach over which the wind can blow, the more energy is absorbed and the greater the wave energy available. As a consequence, the best wave regimes are found where prevailing winds can blow across long stretches of open sea. Once created, waves will travel long distances, and they contain more energy in deep water than in shallow water where they shed energy as they approach the shore.

The largest waves are usually found between 30° and 60° of latitude (hence the name 'the roaring forties' for the region in southern hemisphere where westerly gales blow throughout the year), but there is also a good wave regime in some regions between 30° of latitude and the equator and also in southern polar waters. Regionally, the best locations for wave power exploitation are found on western coastlines along North and South America, in western Europe and western Africa, and the western coasts of Australia and New Zealand.

OTEC energy is found where seas absorb the heat from the sun. As a consequence, the greatest resource is in tropical regions in the seas close to the equator. Much of this thermal energy is far out to sea and difficult to exploit except with floating plants. Shoreline exploitation is possible, but to use the heat in coastal surface waters, an OTEC plant must also have access to deep, cold waters. A depth of a 1 km is generally needed to provide an adequate temperature difference between the hot and cold sources to drive an OTEC heat engine economically. Shore sites with deep offshore waters close enough to the shore to be accessible are rare.

To generate power from the mixing of fresh and salt water it is necessary to have a source of both close together. This situation can generally only be found at the mouth of a river. In principle, high-salinity seas such as the Dead Sea could also provide energy in this way, but exploitation would involve bringing in fresh water and diluting the saline source, which might have significant environmental consequences.

OCEAN THERMAL ENERGY CONVERSION

OTEC takes advantage of the temperature difference between the sea surface and the deep sea or ocean to drive a thermodynamic heat engine that can generate electricity. The best regions for OTEC are usually found within 20° of the equator. Here the surface water temperature can rise to between 25 °C and 33 °C. The hot surface water does not mix with deeper water, so there is a significant temperature gradient between the surface and lower strata. The latter can remain at a temperature of only 9 °C at a depth of 500 m and below 5 °C at a depth of 1 km.

The amount of energy absorbed by tropical oceans is enormous. Estimates from the U.S. National Renewable Energy Laboratory suggest that across the 60 million km^2 of tropical ocean an amount of energy equivalent to as much as 250 billion barrels of oil/day is absorbed. If only a small fraction of this could be converted into electrical power it would provide a massive resource.

An OTEC plant requires a temperature gradient of 20 °C to be able to operate effectively. This means that a surface temperature of 25 °C will normally be necessary and access to colder waters at 5 °C or less. To access this cold water, an OTEC plant must be able to reach at least to 500 m below the sea surface, and for efficient operation it will normally need to pump water from as much as 1000 m down.

A land-based plant, even with deep water close to the shore, is likely to require a pipe of at least 2 km in length to reach sufficiently cold water. The alternative is to build the OTEC plant on a floating platform that is tethered offshore with the cold water pipe hanging from the support, and then carrying the power to the shore via a seabed cable (Figure 14.1). Even with this arrangement, the platform may have to support a pipe of 1 km in length. More generally, if OTEC is to provide anything more than very small amounts of power, plants will need to be able to operate in the open sea where they will have to support a pipe capable of pumping water from up to 1 km below the surface.

OTEC plants not only produce electricity, some configurations can also generate drinking water, which can improve their economics, particularly for small island communities that have little local fresh water. In addition, the cold water that the OTEC plant draws from the ocean depths is rich in nutrients so it

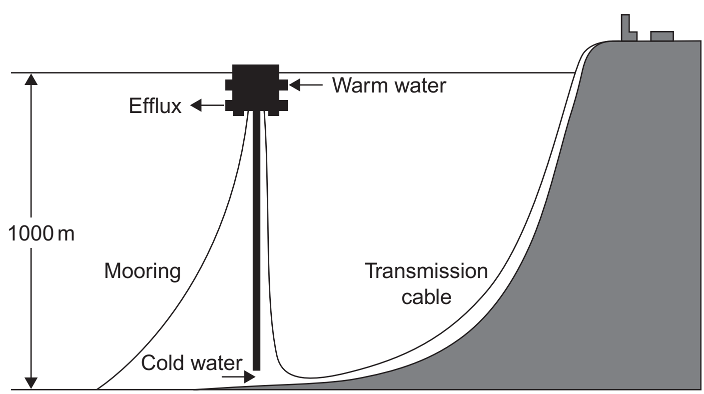

FIGURE 14.1 Schematic diagram of a floating OTEC plant.

can be used to provide an environment for aquaculture. The combination of electricity generation, potable water production, and aquaculture could make multipurpose OTEC systems economically viable in the future.

The OTEC cycle was first proposed by in 1881 Jacques-Arsène d'Arsonval, a French physicist and physician, although the idea was earlier outlined by the author Jules Verne in his book *Twenty Thousand Leagues Under the Sea*. D'Arsonval based his system on a closed-cycle ammonia turbine but never realized his scheme. It was not until 1979 that a project of the same type was built in Hawaii with a net output of 18 kW. Before this, in 1930, a student of d'Arsonval called Georges Claude built a 22 kW OTEC plant based on an alternative open turbine cycle. This plant consumed more energy than it produced and was not commercially successful.

OTEC was revived during the 1970s but the economics have remained difficult. One of the problems is the vast quantity of cold water that must be pumped from a great depth below the surface of the sea. This requires a lot of energy. The volume of water involved for a 100 MW plant is around 200 m^3/sec. To deliver this would require a pipe of around 11 m in diameter and as much as 1 km long. Meanwhile, the discharge of both hot and cold water from a 100 MW plant would be roughly 600 m^3/sec, one-sixth of the flow from the Nile River into the sea.

OTEC Technology

OTEC relies on a very small thermal gradient to drive a heat engine and generate electrical power. The efficiency of a heat engine increases the larger the temperature difference between the hot and cold sources. For a temperature difference of just 20 °C the best theoretical efficiency is 6.6%, rising to 9.6% with a 30 °C temperature difference. Practical efficiencies are likely to be lower than this, and when the power required to pump water from the ocean depths is taken into account, the overall efficiency may fall to 4% or less. As an example, a practical system designed by Lockheed Martin operating with a temperature difference of 18.5 °C recorded an efficiency of 2.65%.

There are two primary cycle types that can be used to exploit this small temperature difference; a closed cycle and an open cycle. The closed cycle system uses a special thermodynamic fluid that is alternately vaporized and condensed as it cycles within a closed turbine system. The most common working fluids for this type of cycle are either a low–boiling point organic fluid that can be easily vaporized using hot seawater, or an ammonia or ammonia-and-water mixture. An organic working fluid will generally be exploited using a Rankine cycle turbine system while the ammonia–water working fluid is commonly found in a Kalina turbine cycle. Both are capable of generating power from small temperature differences.

In a closed cycle system the pressurized working fluid is heated and vaporized by passing it through a heat exchanger through which the hot sea water passes (Figure 14.2). The vapor produced in the heat exchanger is used to drive

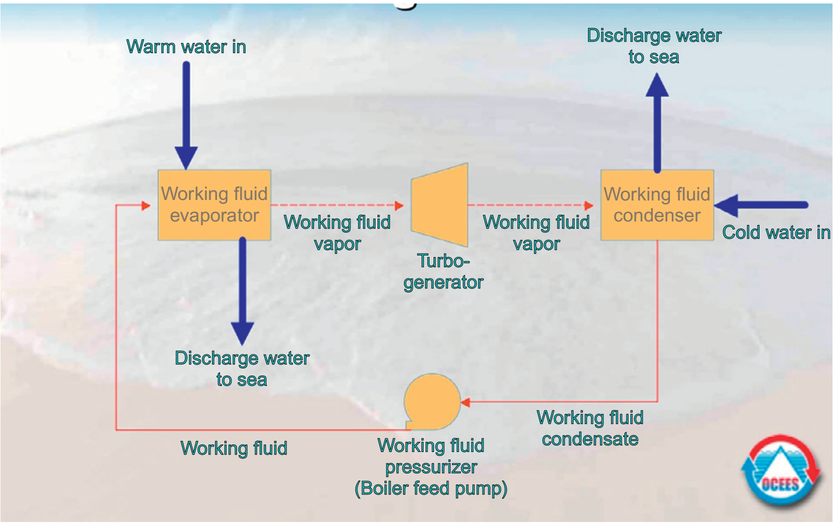

FIGURE 14.2 Schematic of a closed cycle OTEC plant.

a turbine. The vapor exiting the turbine is then condensed using cold seawater in a second heat exchanger before recycling through the first heat exchanger. The difference in vapor pressure between the hot and cold heat exchangers provides the pressure difference that drives the turbine. The turbine then drives a generator to produce electrical power.

The heat exchangers are the most critical components of a closed cycle OTEC plant and they determine if it can operate economically. For high efficiency, the heat exchangers must be very large; a pilot 1 MW plant tested in India in the early part of the first decade of the 21st century included an evaporator, or hot water heat exchanger, with an area of around 4000 m^2 that accounted for 27% of the cost of the plant. As a consequence of the large size necessary, economies of scale demand that such plants must be large to be economical. Closed cycle OTEC plants are unlikely to be viable at sizes less than 40 MW.

The alternative, open cycle OTEC plant does not use a special thermodynamic working fluid but instead uses seawater itself (Figure 14.3). Warm seawater is injected into an enclosure called a flash evaporation chamber that is held at low pressure. Inside the flash chamber a part of the seawater vaporizes and this vapor (steam) is fed into the turbine. After passing through the turbine, the vapor is then condensed in a heat exchanger using cold seawater, creating a pressure gradient that provides the power to drive the turbine.

Flashing seawater in this way produces a vapor between one and three times atmospheric pressure—very low by steam turbine standards. Exploiting this requires a very large steam turbine. Meanwhile, the low pressure in the flash chamber is generated hydrostatically by siting the chamber above sea level so that a head of water to the sea can create a negative pressure.

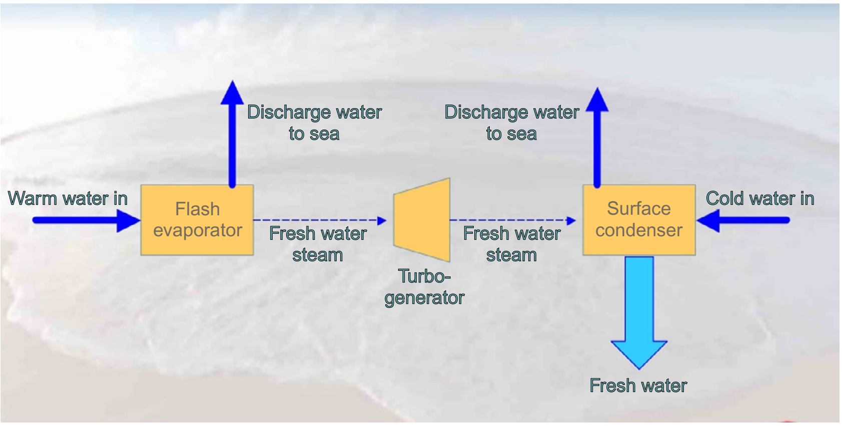

FIGURE 14.3 Schematic of an open cycle OTEC plant.

It is difficult to build an open cycle OTEC turbine with a capacity greater than 3 MW using conventional materials because of the size of the turbine involved. Higher capacities might be achieved with lightweight materials such as those used for wind turbines, and turbines with diameters of 100 m have been proposed. However, it is not clear today that the economics can support the research necessary to develop such turbines. One of the key attractions of an open cycle OTEC plant is that the condenser provides fresh water that can be used for drinking. This might make a plant economical under some circumstances.

It is possible to combine the advantages of open cycle and closed cycle plants in a hybrid plant. Here warm seawater is flashed to produce vapor but the warm vapor then vaporizes, in turn, the working fluid in a closed cycle OTEC system. Both vapors are then condensed using cold seawater. As with a conventional open cycle plant, the hybrid plant can provide drinking water from the water vapor.

While OTEC plants would ideally be built on land, suitable sites are scarce and most of the resource is located far from land in the tropical oceans. To extract this energy requires floating OTEC plants that "graze" the oceans to generate power. Since they would be located far from land, they would then have to store the energy they produced. This might be achieved by converting the electricity into hydrogen via electrolysis of seawater. The approach could become attractive if a wider hydrogen economy develops, but today it is difficult to make an economic case for such a project.

Since the beginning of the 1990s a number of close-shore commercial projects have been proposed, the latest being a scheme to build a 10 MW offshore plant to power a low-carbon resort off the coast of China. This plant would be constructed by Lockheed Martin, which has been developing an offshore OTEC design for several years. However, at the beginning of 2013 no commercial plant had actually ever been built.

One of the main problems with OTEC is environmental. The plant must draw enormous quantities of cold water from the ocean depths, but if this is returned to the surface, it will result in the mixing of sea strata, a cooling of the surface water, and a big increase in the level of nutrients. The consequences of this for local marine life have yet to be assessed. The alternative—returning the water to the same depth as it was extracted—will affect plant economic viability since it will require more pumping, an energy-intensive activity.

There is also the danger that water from the deep ocean will release carbon dioxide dissolved within it when it is brought to the surface. Against this, it has been argued that this nutrient-rich water will encourage marine flora and that this will absorb more carbon dioxide. Only large-scale development will enable these questions to be answered definitively.

WAVE POWER

The waves found in seas and oceans are created when the sea absorbs energy from the wind. The stronger the wind and the longer the reach of sea over which it has to blow, the greater the amount of energy that is absorbed. The rotation of Earth means that the best wind regimes are generally found on western shorelines exposed to the wide ocean and the strongest winds are between 30° and 60° degrees of latitude. Good sites can be found along the western coasts of Europe, North and South America, Africa, Indonesia, Malaysia, New Zealand, and Australia.

The energy contained within waves is manifested as an oscillatory motion of the sea surface. Over very long reaches this can assume a regular frequency, as characterized by the swell found on the oceans, but often and particularly near the coast it will become a superposition of a number of different frequencies. Whatever its precise nature, the motion from the point of view of energy capture is an oscillation of the water surface relative to a fixed point on land or on the seabed. It is this motion that must be exploited in a wave energy converter. (Some converters also exploit the relative motion of two adjacent points on the surface.)

Wave energy is characterized by the amount of energy contained within a 1 m wave front as it arrives at the point of capture. Far out to sea, waves can contain as much as 100 kW/m, and this energy is retained while the waves remain in deep water but the energy content generally falls as the waves approach the coast. At a depth of 20 m, the energy content will typically have fallen by two-thirds.

Average annual wave energy levels along the coast of western Europe are 48–70 kW/m, and in the Americas they can range from 13–102 kW/m with the highest energy levels generally found in the more northerly or southerly parts of those continents. Energy content varies seasonally too, with more energy generally found in winter than in summer.

The possibility of harnessing waves as a power source was first explored during the 19th century, initially as a means of ship propulsion but later as a land-based means of energy capture. The earliest actual wave power machines were built in California between 1890 and 1910 when several "wave motors" were constructed though most were unsuccessful. After this, interest died until the 1970s when projects were launched in the United Kingdom, Japan, and Norway. None of these schemes appeared to be economically viable but interest grew again during the 1990s, and during the first decade of the 21st century, scores of new projects were launched using a wide range of differing technologies. Most are at an early pilot or demonstration stage but interest remains high.

Wave Power Technology

The quest to exploit the oscillatory motion of waves has given birth to a bewildering range of mechanical devices designed to convert that motion into electricity. Some of these are shore based, and some will only operate far out to sea. The principles upon which they operate vary widely too and there is no easy way to classify them.

Shoreline devices include oscillating water column converters and a range of overtopping devices. Oscillating flap converters are usually built close to the shore too. Offshore devices include floats, point absorbers, and wave pumps, all of which need to be tethered to the seabed to operate. Another series of offshore devices, with names such as snake and duck, use the relative motion of a series of floating elements, one against the other as a source of energy. There are also piezoelectric devices that generate electrical energy through the bending of a special material. These can be both shoreline and offshore energy converters.

Shore and Near-shore Wave Converters

Oscillating Water Columns

One of the simplest and most common methods of capturing energy from wave motion is with an oscillating water column (OWC) device (Figure 14.4). This comprises a tube or chamber that has a lower aperture below sea level and the opposite end above sea level, open to the air. The simplest way of envisaging an OWC is to think of a tube with one end immersed in the sea. As waves pass across the tube, the water level within it will rise and fall, alternately forcing air from the top of the tube or sucking it in. This motion of air can be harnessed to turn a type of wind turbine, generating electricity.

Since the movement of air is cyclic, alternately traveling in and out of the tube as waves pass, a conventional turbine would rotate first one way and then the other, and so it would be unable to generate with a conventional generator. There are two solutions to this. The first is to employ a system of valves and two turbines, one that rotates as air is forced from the tube and the other that rotates when air is drawn in. A more elegant solution is to use a turbine that can rotate in

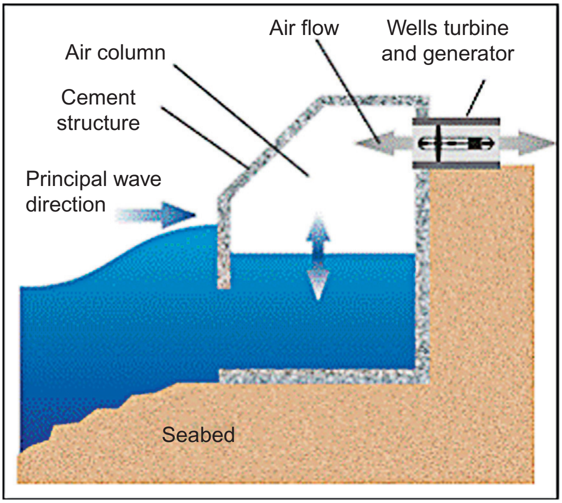

FIGURE 14.4 A schematic of a shore-mounted oscillating water column wave energy converter.

the same sense whichever direction the air flows. The most commonly used of these is a Wells turbine.

OWC devices can either be built on the shore or in shallow near-shore waters. They usually comprise some form of concrete structure anchored to the seabed or shoreline with seawater admitted to an aperture in the part beneath the water level. The device must be large enough so that its lower aperture is always submerged whatever the state of the tide. Otherwise, it will not be able to generate throughout the tidal cycle.

The energy source driving an OWC is variable and conversion into electrical power results in a highly variable output. If the turbine is designed with a significant moment of inertia this can help smooth the output but it will remain highly variable without further smoothing.

A variety of pilot-scale OWC energy converters have been built including a 500 kW unit called the Limpet that was built on the Scottish isle of Islay by a company called Wavegen in 2000. Since then a 4 MW has been proposed involving 30–40 Wells turbines. This project was to be built at Siadar Bay, on the island of Lewis in Scotland, but the project was canceled in 2012.

Overtopping Devices and Tapered Channels

Another simple method for extracting energy from waves is to consider them simply as a variable head of water. If some of the water from the crests of waves can be captured it can be used to create a head of water that rises above the surrounding sea level, and this can then be exploited to drive a hydroturbine and generate power with the water from within the device running back to the sea.

Creating a head of this sort can be achieved in a number of ways. One of the early attempts, called a Tapchannel, funneled waves into a rising and narrowing channel where the kinetic energy of the waves was used to force water to overflow the sides of the channel into a small reservoir built on the seashore. The water was then allowed to run back to the sea through a turbine.

Another approach is to build an offshore reservoir in shallow waters with gently sloping sides so that the energy of the waves as they drive toward the shore forces water to flow over the sides of the reservoir (Figure 14.5). As with the Tapchannel, the water collected in this way is then allowed to flow back to the sea through a turbine.

The head of water created in this way is relatively small and a low-head turbine is required to extract energy from it efficiently. Kaplan turbines have been used successfully in some pilot schemes of this type.

One of the overtopping devices currently being developed is the Wave Dragon. This was originally developed by a Danish company of the same name. Wave Dragon Wales won funding in 2011 to build a 7 MW prototype off the Welsh coast of Pembrokeshire. However, in 2013 the energy converter was still not constructed.

Oscillating Flaps

A third common near-shore wave energy converter is the oscillating flap device, sometimes also called an inverted pendulum converter (Figure 14.6). The basic principle behind this type of converter is to devise a buoyant flap, the bottom of

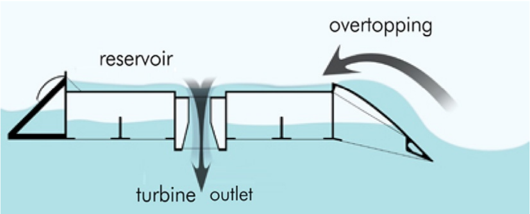

FIGURE 14.5 Schematic view of an overtopping wave energy device.

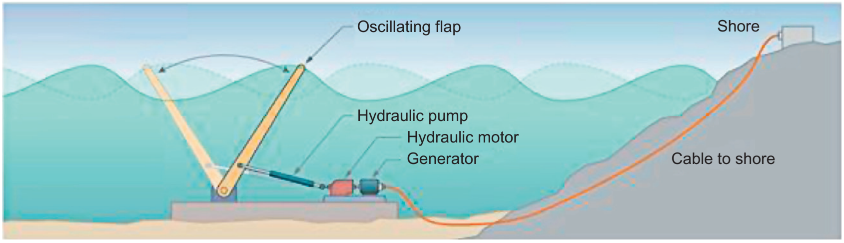

FIGURE 14.6 Oscillating flap wave energy converter.

which is hinged and attached to a foundation anchored to the seabed. The body of the flap then rises under water above the hinge. As waves move across the site of the device, they cause the flap to oscillate backwards and forwards and this oscillatory motion is converted into electrical power.

A similar principle can be exploited using a foundation structure that rises close to the surface of the sea. In this case the structure supports a weighted flap that is hinged from the top and falls down into the sea. Again, as waves cross the site of the device they cause the flap to move backwards and forwards.

Energy conversion with such devices is often carried out by using the oscillating motion to power a pump that drives a fluid, often seawater, under pressure through a pipe to the shore. A high-pressure flow of water of this type can then be converted into electricity using a Pelton turbine, normally used for high-head hydropower plants. With a system of this type, several wave energy converters can be used to provide pressurized water to a single turbine. In addition, water can be stored under hydrostatic pressure in an accumulator, allowing this type of wave energy device to store energy and smooth its output, providing a more valuable source of power for the grid.

One oscillating flap device that has been under development is the Waveroller, from Finnish company AW-Energy. The company deployed a pilot scheme off the Portuguese coast in 2012 comprising three 100 kW Waveroller oscillating units. The company is hoping to be able to use the pilot to develop a 500 kW commercial unit. A second oscillating flap device called the Oyster has been developed by U.K. company Aquamarine Power. A single 800 kW demonstration unit was installed at the European Marine Energy Centre, Orkney, in 2012. This uses the pump principle to bring pressurized water ashore to drive a hydroturbine.

Offshore Devices

Buoyancy-based Devices

One of the principal categories of offshore wave energy converters is based on a floating device such as a buoy that is tethered to the seabed. While such devices can operate close to shore, since the energy contained in waves is greater farther offshore they are better deployed in deep water.

A buoyancy device will be attached to an anchor that is fixed to the seabed. The device itself will then float to the surface, or it might be held just below the surface. In either case, the passage of waves will create an oscillating force acting on the buoy as passing waves pull it toward the surface and away from the seabed. This force is then used to generate electricity.

The force can be converted into a useful form of energy in a number of ways. One method involves installing a flexible hose as part of the tethering anchor. This hose is alternately stretched and released as waves pass, applying a pressure cycle to the tether. The motion causes the tube to contract and expand,

and this can be used to generate a hydrostatic pressure that can be converted into electric power. Another, similar, approach is to use a piston pump fitted to the tether or buoy. The movement of the buoy up and down drives a piston in and out of a cylinder integrated into the device and this can be used to generate hydrostatic pressure too.

The main alternative to using the motion to pump water is a linear generator. This is a device that can produce electricity from linear rather than rotary motion (Figure 14.7). If a generator of this type is fitted into the tether of a buoy, the up and down motion as the buoy is lifted and dropped by waves can then be turned directly into electrical power.

Pontoons, Snakes, and Ducks

The second principal category of offshore wave energy converters comprises a disparate group of machines that have one thing in common. Each is made up of a number of floating sections that are joined end to end through a hydraulic linkage. When one of these devices is taken out to sea the sections will move relative to each other and this relative motion is used to extract power.

One of the earliest wave power devices, called the Salter's Duck after its inventor, Stephen Salter, was a converter of this type. Devised in the 1970s, it comprised a large anchored section to which was fitted a lighter hinged "beak." When moored offshore, wave motion caused the beak to oscillate up and down and this motion was exploited to provide power.

Some devices of this sort are made up of two or more elements of the same size. The relative motion of these as they jostle under the influence of the waves

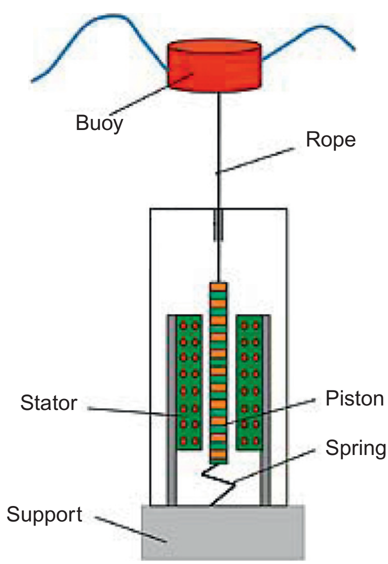

FIGURE 14.7 Floating wave energy converter with linear generator.

at sea then allows power to be extracted. Others are designed with one large, relatively immobile element and another that is tuned to respond to the motion of the waves. This latter group is often called point absorbers and they may be tethered to the sea bed, like buoyancy devices, above. Tuning a device in this way will normally mean that it only responds to a band of wave frequencies and is unaffected by others outside this band. Whatever the design, however, the principle is the same: one part of the device moving relative to the other, as shown in Figure 14.8.

As with other offshore devices, energy can be extracted by using the motion to pump water, by generating hydrostatic pressure that can be used to produce electricity, or by using some sort of linear generator. An example of this latter approach is the Archimedes Wave Swing, an underwater tethered buoy that was originally developed in The Netherlands but is now owned by U.K. company AWS Ocean Energy. A 2 MW prototype was tested off the Portuguese coast in 2005 where it achieved an output of 1 MW. More recently, AWS Ocean Energy has been awarded £3.9 m to build its latest prototype.

Another tethered buoy is the Powerbuoy being developed by U.S. company Ocean Power Technologies. Like the Archimedes Wave Swing this is a point absorber tuned to extract energy from ocean waves. The company has tested several prototypes but is currently developing what is proposed to be its first commercial deployment off the coast of Oregon. This will initially involve one 150 kW unit followed by 10 grid-connected units providing a total of 1.5 MW. Following this, the company plans a 50 MW installation. Factory testing of the first unit was completed in mid-2012.

Yet another tethered buoy is the Aquabuoy, in this case a floating device, which has also been tested off the coast of Oregon. There were plans for a 2 MW grid-connected power plant selling power to a U.S. utility but there have been no recent reports of progress.

Of the second type of offshore floating devices, the main example today is the Pelamis developed by Pelamis Wave Power. This comprises a series of buoyant cylindrical sections linked through hydraulic joints. Movement of the sections pumps oil that is used to drive a generator. A 750 kW prototype comprising three 250 kW units was deployed off the coast of Scotland in 2005, and in 2008 three units were installed off the Portuguese coast but had to be brought back ashore later that year. Since then there appear to have been no further units installed.

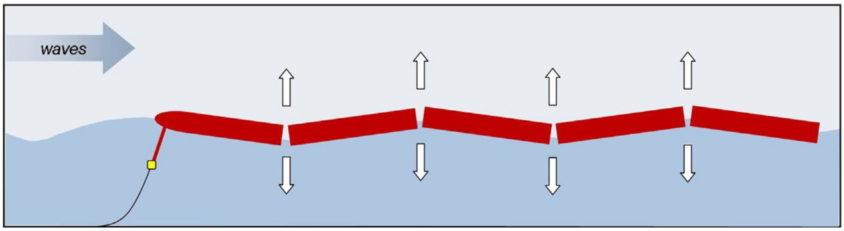

FIGURE 14.8 The action of the waves on the hinged sections of a floating wave energy converter.

Piezoelectric Devices

Piezoelectric materials produce an electric voltage when they are placed under stress. This may be a result of bending, stretching, or a variety of other means including when they are subjected to vibrations. Any and all of these means of generating electricity could potentially be exploited in a wave energy device.

One idea that has been put forward involves piezoelectric "seaweed." Wafers of a piezoelectric material are anchored to the seabed so that the wash of waves across them causes them to bend backwards and forwards, generating electricity. A second involves the stretching of a piezoelectric material that could, in principle, be used to generate power from a tethered buoy if the generator was mounted into the tether. However, while there are a wealth of possible ways such materials could be used, no wave converter using them has yet been built.

MARINE CURRENT ENERGY

Marine current converters, sometimes also called tidal stream converters, utilize the energy contained in flowing water generated by the motion of the world's tides as a source of electrical power. Currents of this type can be found in many estuaries around the world and also in coastal straights and areas where there are offshore coastal features that allow the tidal movement to create a significant current. In principle, the same type of converter can be used in a river, too, as an alternative to a conventional hydropower plant.

The amount of energy that is available in a marine current depends on both the amount of water that is moving and the speed at which it flows. Areas where there is a large tidal reach are likely to offer the best marine current potential since the volume of water moving on each tide will be large. Meanwhile, the faster that water flows, the greater the energy it contains.

The energy that flowing water contains varies as the third power of its velocity so energy content is extremely sensitive to speed. A site with a fast current will therefore offer the most efficient way of harnessing this source of energy. Such sites may be found where the natural flow is constricted either in a river or offshore. Average flow rates of 2 m/s or more are generally considered ideal for a marine current power plant.

While some marine current plants may take energy from flowing river water, most will exploit tidal motion. The latter is both cyclical and variable, with the strength of the current depending on the state of the tide. This needs to be taken into account when designing a marine current project. There will be a small seasonal change too as the tides vary between spring and neap tides (the two extremes when the gravitational pulls of the moon and the sun either reinforce or oppose one another). The flow in river sites, on the other hand, will have little short-term variability but may vary significantly from summer to winter.

Flowing water is conceptually similar to flowing air. In consequence, the most common methods of extracting energy are similar to those developed for wind power—that is, turbines with several blades that rotate in the flow. However, because water is much denser than air, the energy density of flowing water is much greater so turbines can be much smaller for the same power output. For example, a 10 m diameter turbine in water flowing at 4 m/s can generate around 880 kW, whereas a 10 m wind turbine turning in a wind speed of 4 m/s would barely achieve 0.5 kW. For a similar output, the wind turbine would need to have a diameter of around 70 m.

In a 2 m/s flow of water an 18 m diameter turbine is needed to sweep out (or intersect) water carrying a power density of 1 MW. At 4 m/s the turbine size needed to sweep out a power density of 1 MW falls to 6 m. Only around one-third of the energy swept out is captured and converted into electricity, so for a 4 m/s flow a 1 MW marine generator would need to be a little more than 10 m in diameter.

Marine Current Energy Converters

Based on the analogy with moving air, most marine current converters are turbines that rotate in the water flow, providing a mechanical power output that can be converted into electricity. Many exploit technology similar to that used by the wind industry. The conventional wind turbine in use today is a horizontal-axis wind turbine mounted on the top of a tall tower. Exactly the same type of technology can be used for marine currents with, in this case, the tower fixed to the sea or riverbed.

It is possible to carry this analogy too far and there are often significant differences between wind and water turbines, particularly concerning the means of mounting the latter in the water flow. While all wind turbines are placed on towers, marine current turbines can be mounted on towers, they can be hung from the bottom of floating supports, and they can be mounted on buoyant structures that are anchored to the seabed or riverbed. Generators may vary too with some marine devices using rim generators that allow greater flow through the turbine.

There is also a viable alternative to the horizontal-axis wind turbine for marine power: the vertical-axis turbine. Various forms of the latter were tested during the early days of wind power development, but as the technology matured the designs were mostly abandoned for the standard design in use today. The vertical-axis turbine has several significant advantages for marine current deployment. The most important of these is the ability to continue to rotate in the same direction, whichever way the current flows. The vertical shaft of the turbine, meanwhile, allows the generator to be mounted at one end. This might either be at the surface if the turbine is deployed from a floating barge, or on the seabed.

Another way of extracting power from a tidal flow is with a paddled wheel, conceptually identical to the old mill wheel used to drive water mills. Such wheels or turbines are only partly submerged, with water flowing against the submerged part of each paddle. Technically, this is known as a cross-flow turbine. Such devices are being used today but they are generally not as efficient as fully submerged turbines. However, the technology is simple, which may be advantageous for small installations. (There are fully submerged versions of cross-flow turbines available too.)

There are a number of other means of extracting power from flowing water. These include oscillating airfoils that rise and fall in the flow, with the oscillating movement used to drive a pump or generator, and devices based on the Venturi effect when water is constricted to flow through a narrow passage.

Horizontal-axis Turbines

A horizontal-axis marine turbine will often share many features with a similar wind turbine. A typical device may have a three-bladed rotor mounted on a shaft that drives a generator through a gearbox. Since it will be operating under water, all the electrical and mechanical components must be protected in a watertight container. Gearboxes have been notorious as a point of weakness in wind turbines and designers have developed direct-drive systems that eliminate them. Similar systems can be applied to marine turbines. Variable-speed generators have also been deployed in the wind industry, with electronic AC-DC-AC converters to match the output to grid frequency. Again these are becoming common within the marine industry. See Figure 14.9 for a schematic of a horizontal-axis turbine.

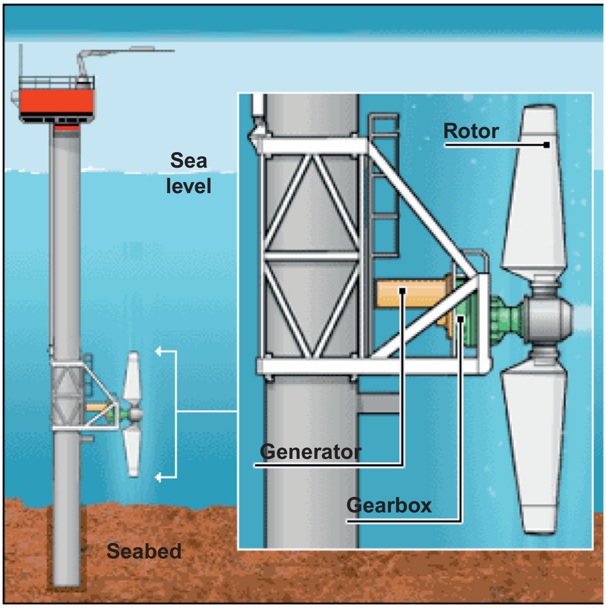

FIGURE 14.9 Schematic of a horizontal-axis turbine, marine current converter.

While many tidal stream turbines have three blades some may have two and others more. These may be of a fixed pitch, which is structurally the simplest design, or they may be capable of varying the pitch to control rotational speed. Speed control is generally less critical for tidal stream turbines since the water speeds are much more predictable than wind speeds. In addition, whereas wind turbines require the ability to be shut down in high winds, this facility is not needed in a marine turbine since such extremes do not generally exist (although exceptionally, they might be encountered on a flooded river).

Many marine turbines will have open blades but some designs include a shroud that encloses the blades. This shroud can serve two purposes: to protect the blades from damage and prevent marine animals from being damaged by the blades, and to control the water flow.

A shroud can make the turbine more efficient in a cross-flowing current, which is likely to be common in the marine environment where current directions change with the tide. A correctly designed shroud can also improve efficiency by acting as a diffuser that flares on the downstream side of the turbine. This serves to increase the flow through the turbine and can increase efficiency from 35% for an open turbine to 60% for a well-designed shrouded design.

Marine turbines with shrouds lend themselves to an unconventional type of generator design called a rim generator. In this type of machine the generator rotor has no central shaft but is instead mounted on the blade tips of the rotor that penetrate into the surrounding shroud. The rotor itself forms a part of the generator. The generator stator, meanwhile, is built into the shroud. Such a design has only one moving part—the rotor—and since there is no axially mounted generator, there is less impediment to flow through the rotor, leading to higher efficiency. However, the technology is less well tested than conventional generator design and remains to be proven.

There are several means of mounting horizontal-axis wind turbines in the water flow. The simplest, conceptually, is a monopole tower that is sunk into the seabed or riverbed with the turbine mounted on top of it. However, these can be expensive to install. An alternative is a gravity foundation that holds itself in position by means of its mass.

Buoyancy can be utilized too. One approach is to make the turbine structure buoyant and anchor it to the seabed. Such tethering systems are not as stable as a monopole but they are much cheaper. Turbines can also be hung from floating platforms.

Ease of maintenance is always an issue in a marine environment. Since maintenance cannot be carried out with ease under water, a means must be provided to raise the turbine out of the water. Some monopole structures are telescopic and can be raised to bring the turbine to the surface. In other cases, the monopole extends above sea level and the turbine can be raised and lowered on it. Buoyant turbines anchored to the seabed can be raised by untethering them, while those on floating platforms can be raised onto the platform for maintenance to be carried out.

All marine mounting systems are expensive and it is often cost effective to mount two or more turbines on a single support. They are often mounted in pairs, with each member of the pair rotating in the opposite direction to the other to balance the rotational forces about the mount.

Vertical-axis Turbines

A vertical-axis turbine has its blades mounted onto a vertical shaft (Figure 14.10). The vertical-axis design has two significant advantages: it will rotate whatever direction the water flows, and its generator can be mounted at the top or bottom of the shaft.

The insensitivity to flow direction means that a vertical-axis turbine will operate efficiently in a tidal region where the flow reverses twice a day. It will also be unaffected by cross-flows that might reduce the efficiency of a horizontal-axis turbine. Depending on the site this could prove a great advantage. Meanwhile, the ability to mount the generator at either end of the shaft means that it can be placed either on the seabed or riverbed or, more usefully, it can be deployed above sea level either from a shaft extending beyond the sea or on a floating support.

Vertical-axis turbines can come in a variety of shapes. One of the most popular from the early days of wind turbine development is the Darrieus or egg-beater design, so-called because its curved and twisted blades resemble the kitchen implement. Others have vertical blades mounted onto the shaft using horizontal struts to create an H-shaped rotor.

The H-shaped rotor has an additional advantage—it sweeps out a rectangular area as opposed to the circular sweep of a horizontal-axis turbine. This means the dimensions of the rotor can be tailored to the shape of the

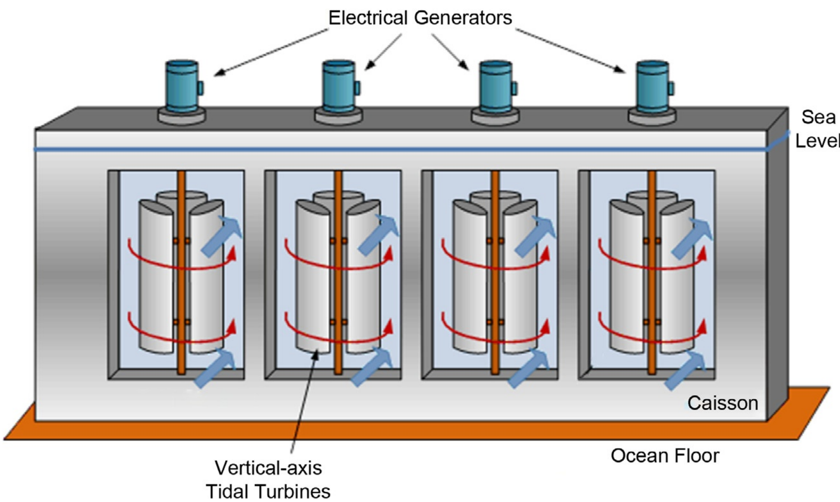

FIGURE 14.10 An array of vertical-axis marine current converters.

waterway into which it is to be sited. Arrays of vertical-axis turbines can, in principle, more efficiently collect energy from a water flow since they can be arranged to intersect a larger cross-sectional area.

Some vertical-axis turbines will have only two blades, others more. It is also possible to build a vertical-axis turbine with a rim generator. This type of machine, again, has no central axis. As a further variation, there is at least one design of vertical-axis machine that provides hydraulic power to the shore instead of incorporating a generator.

Cross-flow Turbines

A turbine that is like an old-fashioned water wheel can be used to generate power from all types of flowing water. Tidal mills from earlier centuries used this type of technology and modern versions can be adapted to a range of sites. This type of turbine is an efficient energy converter, but since only a part of it captures energy from the flowing water, it requires a larger turbine for the same power output as a fully submerged horizontal- or vertical-axis machine. The turbine has the additional disadvantage that it cannot operate on both the flow and ebb of a tidal site.

There is a variety of cross-flow turbines that will operate fully submerged. The turbine is similar to the anemometer used to measure wind speed. Blades are cup-shaped so that from one direction they capture the water but from the other they present a much lower water resistance. Though less efficient than more conventional types of turbines, these can be effective in small installations.

Other Marine Current Devices

In addition to the variety of turbines, there are other more unconventional marine current devices that have been developed and tested. One of these is based on a hydrofoil. A hydrofoil is a shaped blade that is intended to produce lift in exactly the same way as an airplane wing. To provide energy conversion, a section of hydrofoil is attached to the end of a beam, the opposite end of which is attached to a pivot. When this is placed in a flowing current with the hydrofoil horizontal, close to the bed of the river or sea, the flow of water will create lift, forcing it upwards and causing the pivoting beam to rise. However, when the angle of attack becomes too great, the hydrofoil instead generates drag and this forces the beam down again. See Figure 14.11.

The oscillating motion will continue indefinitely so long as water continues to flow. The cyclic motion of the device can be used to drive a piston in and out of a cylinder and this can be used in a number of ways to produce electrical power. For example, it can be used to provide hydraulic pressure that can be utilized ashore, or it can be converted into rotary motion (just as in a piston

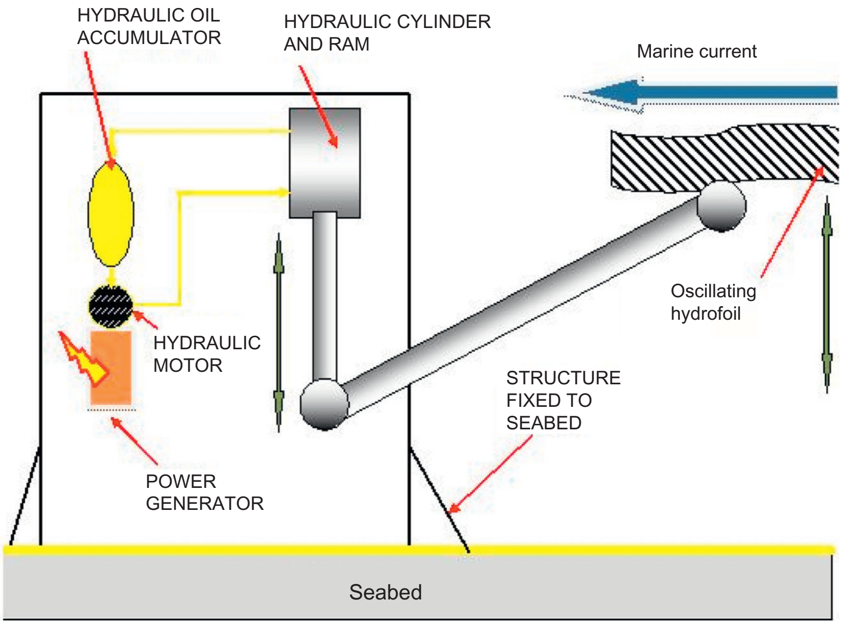

FIGURE 14.11 Schematic of an oscillating hydrofoil marine current converter with hydraulic power transmission.

engine) to drive a generator. The hydraulic version can also store energy in a hydraulic accumulator.

A variation on the basic hydrofoil design involves mounting the hydrofoil vertically at the end of the beam instead of horizontally. This generates an oscillatory motion from side to side, which can be exploited similarly to the vertical device.

The Venturi effect, found when a fluid is forced through a constricted opening, has also been exploited to generate power from marine flows. The Venturi effect is a pressure drop found within the fluid as it passes through the constriction, and in most cases this is used to enhance flow and increase efficiency of a turbine. However, one device uses the pressure drop to provide suction, generating an air flow that is used to drive a wind turbine.

Another novel system for generating power from marine flow proposes to utilize the magnetohydrodynamic effect, whereby an electric current is generated when a fluid that conducts electricity is passed through a powerful magnet. In a system put forward by Netherlands company Nepture Systems, salt water, which is an electrical conductor, is passed through a powerful magnetic field and electrical power is extracted. The system might be combined with a shoreline superconducting magnetic energy storage device that could both store energy from the generator and provide the powerful magnetic field needed for its operation.

Marine Current Projects

A number of marine current designs have reached the pilot stage of their development. One is Seagen, a device developed by U.K. company Marine Current Systems that now has Siemens as its major shareholder. The company has tested a pilot generator in Strangford lough in Northern Ireland. The demonstration unit comprised two 16 m horizontal-axis turbine trains, one mounted on each end of a horizontal beam that is attached to a vertical support. The power trains can be raised and lowered on the vertical support for maintenance. The pilot scheme had a generating capacity of 1200 kW. The company is now building a commercial 2 MW unit with two 20 m rotors.

Another U.K.-designed device is Deltastream, devised by Tidal Energy Ltd. This comprises a tripod structure that sits on the seabed with a turbine at each corner. Initial generating capacity is likely to be around 1 MW. The device is due to be tested in U.K. waters at the beginning of 2014. Meanwhile, Irish company OpenHydro has developed a shrouded turbine with a gravity base that it designed to sit on the seabed. This machine also has a rim generator and the turbine has a hole in its center to allow marine creatures to pass. The company has tested a 6 m pilot turbine.

In the United States, a company called UEK has developed the Underwater Electric Kite, essentially a device that comprises a pair of buoyant shrouded turbines, mounted together, tethered to the seabed and the whole structure flown like a kite in the flowing current. The design has been tested at a small scale but no major demonstration project appears to have been completed.

Vertical-axis turbines are the basis for Blue Energy Canada's tidal fence concept. This involves deploying an array of vertical-axis turbines in a fence or barrage across a tidal estuary. A vertical-axis turbine is also the basis for Canadian company New Energy Corp.'s EnCurrent power generation system.

The Stingray is a hydrofoil-based oscillating generator. A 150 kW prototype was tested off the Scottish coast where it developed 40 kW in a 2 m/s current, but development appears to have halted.

SALINITY GRADIENT POWER GENERATION

Salinity gradient power generation depends on the natural process of diffusion, or mixing. To understand how this can be used to generate electrical power, imagine a special box, divided into two compartments, as shown in Figure 14.12. One of the compartments is filled with salt water and the other with fresh water. If the division between the two compartments were carefully removed, without creating any turbulence, within a short period of time the water in both compartments would be salty. Diffusion would have mixed the fresh and salt water.

If the division between the two compartments is now replaced by a special, semi-permeable membrane that will allow water to pass through it but will not allow any

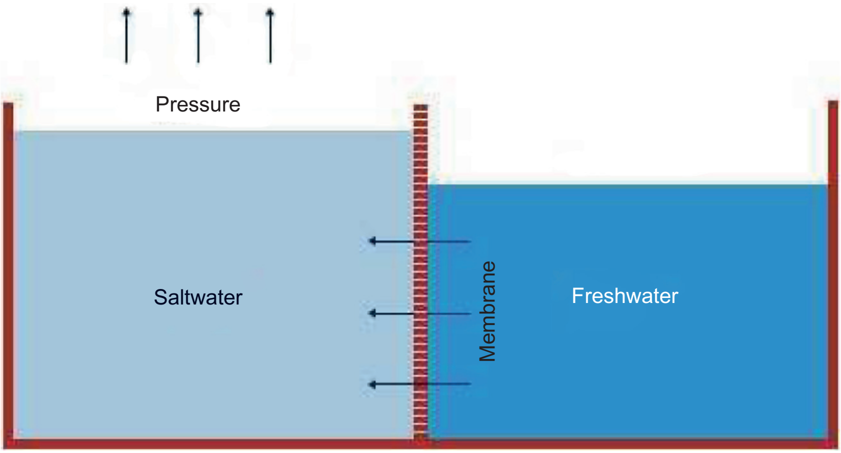

FIGURE 14.12 The principle of salinity gradient power generation based on osmosis.

dissolved salt to cross, when one side is filled with salt water from the sea and the other side with fresh water, water will pass through the membrane from the fresh side to the salt side to try and equalize the concentration of salt in each chamber. Since this is not possible, in theory the water will continue to cross until a head of water builds up on the salt side that can create a pressure sufficient to prevent further water crossing. For seawater, this equilibrium would require a head of 240 m.

Osmotic Power

The pressure that can be generated as a result of diffusion is clearly significant. In principle, it should be possible to exploit this pressure to generate electric power. This is the idea behind the osmotic power plant, first proposed by Sidney Loeb during the 1970s.

In an osmotic power plant a cell is created containing pressurized seawater and this is opened, through a semi-permeable membrane, to fresh water. Fresh water will then diffuse by osmosis into the pressurized chamber increasing the pressure further. The excess pressure is then released by allowing some of the salt water to flow out through a turbine.

The development of osmotic power has been taken up by the Norwegian utility Statkraft that built a prototype plant that began operating in 2009. The plant has a generating capacity of 2–4 kW, although the company predicts this can be raised to 10 kW with better-performing membranes.

Statkraft is proposing to build a much larger plant in Norway to demonstrate the principle. The proposed size is 1–2 MW. The company has estimated that the global potential for osmotic power generation is around 1600 TWh annually.

Plants exploiting this technology would be sited at places where fresh water flows into the sea and would harness the mixing to generate electricity.

Vapor Compression

The vapor pressure of water that is generated above fresh water and salt water, if both are at the same temperature, is slightly different. Although the pressure difference is not great, it is possible to generate a vapor pressure from fresh water, under partial vacuum conditions, and then condense the vapor in salt water. This will create a pressure gradient that can be used to drive a gas turbine.

The pressure difference available is very small and would require a very large turbine, similar to an open cycle OTEC plant. This system does away with the need for a semi-permeable membrane but may be more difficult to develop than osmotic power.

Hydrocratic Power

If fresh water is introduced into the bottom of a vertical tube, pierced with holes, that is situated in salt water, then the salt water will force its way into the tube—driven by diffusion—to dilute the fresh water. The flow of water upwards through the tube is thus reinforced and a turbine placed in the flow can apparently generate more energy than is required to pump the fresh water into the pipe in the first place.

The process, christened hydrocratic power, has been patented and tests of its validity have been carried out. The scheme might be installed where power plant cooling water is discharged.

Reverse Electrodialysis

In principle, fresh and salt water can be used at two electrodes of a specially designed battery to produce an electric current directly. To achieve this, a series of cells are created with each separated from the next by a semi-permeable membrane. In this case, one membrane will only allow dissolved sodium ions to pass through it while the next will only allow dissolved chloride ions. Salt water is pumped into the first cell, fresh water into the second, salt water into the third, fresh water into the fourth, and so on along the line. Sodium ions will then diffuse from the salt water on one side into each fresh water cell, while chloride ions will diffuse from the other side. This will create a voltage separation between the saltwater cells on either side of the freshwater cell.

This process, called electrodialysis, was observed by R. Platte in 1954. The concept is the basis for the development of a power generation system by Netherland's company REAPower.

COST OF MARINE POWER GENERATION

All the marine power generation technologies in this chapter are at an early stage of development of commercialization and realistic installation costs are difficult to establish. Based on recent estimates, the most expensive is a floating OTEC plant that has an estimated cost between \$4200/kW and \$12,300/kW.[1] Wave power plants are somewhat less expensive with an installation cost between \$4100/kW and \$6300/kW. Tidal stream plants have an estimated cost between \$1700/kW and \$4000/kW. Salinity gradient technologies are too new for any reliable estimates of installed cost to be available.

While all these technologies appear relatively expensive to install, the energy has no cost and so the electricity they produce should be somewhat cheaper than these installation costs indicate. Wave power–generated electricity appears to be the most expensive today with an estimated cost of more than \$180/MWh. For OTEC, the cost estimates are \$70–220/MWh and for tidal stream they are \$40–120/MWh. Cost estimates for osmotic pressure power are \$70–130/MWh.

1. Costs are taken from *The Future of Marine Technologies*, Global Business Insights, 2010.

Chapter 15

Biomass-based Power Generation

Biomass is a fuel derived from plants and animals. The material classified as biomass can include wood and wood derivatives, a variety of agricultural and animal wastes, a part of urban waste, and some industrial wastes and by-products.

In the past biomass was the most important source of energy for virtually all regions of the globe with societies relying almost exclusively on fuels such as wood and charcoal for heating and cooking. The Industrial Revolution introduced new forms of energy including coal, oil, and eventually electricity, and these have taken the place of biomass across the developed world. However, biomass remains an important fuel in underdeveloped regions. According to the World Energy Council's *Survey of Energy Resources*, biomass still supplies 50 EJ of primary energy each year or around 10% of global primary energy consumption. Most of this is used for heating and cooking.

Although biomass can be burned like coal in a furnace to provide heat for steam generation, the use of the fuel for power generation has not yet been widely adopted. Where biomass power generation plants do exist, they are often linked to industries that can supply them with fuel. So, for example, the wood-processing industry, which produces a variety of wood wastes such as sawdust, has a long history of biomass power generation, as does the paper industry. Agricultural wastes can also provide fuel for a local facility. In recent years global warming has raised the profile of biomass as a potential renewable source of energy, and with this interest, its use has begun to expand beyond the use of wastes to include specially grown biomass crops supported by a growing trade in biomass fuel.

Even with this new interest, the size of the biomass power generation industry remains limited by the availability of fuel. Wastes, though commonly available, can only provide a limited global capacity, and if the technology is to expand significantly, large plantations of fuel crops will be required. Such plantations are beginning to appear, but in many cases they compete with agriculture for land normally used for food production and this can lead to opposition. In a world where many people still do not have sufficient food for their everyday needs the environmental credentials of biomass crops have yet to be clearly established.

Meanwhile, global biomass power generation today consists of a relatively small number of power plants that burn biomass exclusively and some coal-fired power plants that burn a small amount of biomass with coal to reduce their net greenhouse gas emissions. Many of these plants of both types burn wastes and most of them are to be found in either the United States or Europe. The actual size of global biomass power-generating capacity is difficult to estimate because many of the plants are small. According to one report, the size of the global capacity of dedicated biomass plants in 2013 was 59 GW[1] or 1.2% of global generating capacity. Around half of this capacity is located in Europe. There are indications that biomass-generating capacity growth is accelerating in response to global warming, but overall capacity remains small compared to solar- and wind-powered generation, the two major renewable resources in use today. Predictions for future growth vary widely but capacity could reach between 80 GW and 120 GW by 2020. Potentially, however, the industry could become much larger if biomass resources that have so far remained untapped were brought into use.

There are a number of different ways of converting biomass into energy. The simplest and most widespread is to burn the fuel in a furnace and use the heat produced to generate steam that drives a steam turbine. Most existing plants of this type are small and relatively inefficient but technologies such as biomass gasification can improve efficiency significantly, as can an increase in plant size. It is also possible to mix a proportion of biomass fuel with coal and burn it in a coal-fired power plant, a process called co-firing.

Another approach is to produce liquid biomass fuels either by fermentation of crops such as sugarcane or from oil-bearing crops such as sunflowers and rape seed. In principle, such fuels can be used in piston engine or gas turbine–based power generation systems. However, most of these liquid fuels are being used for transportation and that is likely to remain their main application in the near future.

Biomass is considered to be greenhouse gas neutral. This is because while the combustion of biomass will generate carbon dioxide just as would the combustion of a fossil fuel, when that biomass is regrown, it will reabsorb the same amount of carbon dioxide as it released during combustion. Thus, the growing, burning, and the regrowing of biomass simply cycles carbon dioxide between the atmosphere and biosphere.

From a renewable energy perspective biomass has a number of advantages compared to other renewable sources. One of the most important is that biomass generation, based on a combustion power plant, does not rely on an intermittent source of energy and so it can be controlled to provide power when it is needed,

1. Market Data: Biomass Power Generation, Navigant Research, 2013.

just like a conventional fossil fuel–fired power plant. Second, because biomass power generation is based on combustion technology, utilities are familiar with it and therefore comfortable with adopting it.

The main problem with biomass is that under most circumstances the economics of power generation are not favorable. Biomass combustion plants have similar costs to conventional fossil fuel–fired plants of the same type, but because they tend to be less efficient, overall production costs are higher. This is not a problem where there is a ready source of cheap fuel available, which is why industries like wood and paper have adopted it widely. When biomass fuel must be purchased the technology is often not competitive, although costs are improving. However, it can become more favorable when there are subsidies for renewable energy production and it is in regions where these are available that the technology is beginning to gain ground. The International Energy Agency has predicted that the use of biomass for electricity generation could double between 2010 and 2020 and that by 2035 it could be four times higher than in 2010, based on current conditions within the electricity industry. With more aggressive promotion of renewable energy, usage could conceivably rise significantly higher.[2]

TYPES OF BIOMASS

The global biomass resource is the vegetation on Earth's surface. This is equivalent to around 220 billion dry tonnes, or 4500 EJ (4500×10^{18} J), of energy according to the World Energy Council. Between one- and two-thirds of this (the proportion depends on the means used to estimate the amount of carbon fixed annually) are regenerated each year by photosynthesis. At the beginning of the 21st century, as already noted, biomass equivalent to 50 EJ was being used each year to provide energy, mostly from wood fuel for heating and cooking. Estimates suggest that between 200 EJ and 500 EJ could eventually be utilized for power generation. With primary energy demand expected to reach between 600 EJ and 1000 EJ by 2050, biomass sources could, in principle at least, provide a significant proportion of total demand.

From the perspective of power generation biomass can be divided into two categories: biomass wastes and energy crops. Biomass wastes are the most readily available forms of biomass but their quantities are limited. Energy crops, grown on dedicated plantations, are more expensive than wastes but they are capable of being produced in much larger quantities as and where required. Location is important because biomass has a lower energy content than coal and cannot be transported cost effectively over great distances. It is normally considered to be uneconomical to transport it more than around 100 km, although some biomass is traded over much greater distances today, confounding

2. *World Energy Outlook 2012*, International Energy Agency.

conventional economics. If energy plantations are established close to a biomass power plant, transportation costs can be minimized.

BIOMASS WASTES

Biomass wastes can be divided into four categories: urban waste, agricultural waste, livestock waste, and wood waste. Urban biomass waste is a special category, available in relatively small quantities. It usually comprises timber waste from construction sites and some organic household refuse together with wood and other material from urban gardens. Most of this is cycled through an urban refuse collection and processing infrastructure where the biomass waste must be separated from the other refuse if it is to be burned as fuel. (However, more urban waste can be burned in a specialized power-from-waste plant and this may make separation of biomass waste from household refuse uneconomical.) While separation is an expensive process there is often a fee available for disposing of the waste and this helps keep fuel costs low.

Some organic urban waste ends up in landfill sites, although the use of landfills is being reduced in most developed regions because of the growing need to recycle. When organic waste is buried in this way it can generate methane through anaerobic decomposition underground. This methane is both an extremely potent greenhouse gas and a potential hazard, but it can be collected relatively easily and then either flared or used in a power generation plant to provide electricity.

Agricultural wastes, often referred to as agricultural residues, are one of the most important sources of biomass today. These are available throughout the world and they include a number of very important biomass resources. Across Europe and North America there are enormous quantities of wheat and maize straw produced each year. These farming residues are valuable fuels but they are seasonal and therefore require storing if they are to provide a year-round supply for a power station. Sugarcane processing produces a waste called bagasse at the processing plant where it can easily be utilized to generate electricity. Further waste, called trash, is left in the field but this can be collected and used too. Rice produces straw in the fields and husks during processing—both potential fuels. The shells and husks from coconuts can be used to generate electricity as can waste from oil palms, while the periodic recycling of oil palms and rubber trees (plantation trees have a life of 20–30 years) can provide wood waste for power generation. Many other crops from all parts of the world produce stalk waste that can be utilized. Indeed, wherever crops are grown and harvested there is normally some residual material that can be used as a source of energy.

There is one important caveat to the use of agricultural wastes as power plant fuel. From the perspective of sustainability it is important that some biomass material is returned to the soil after a crop has been harvested if the soil is to retain its fertility. If all the biomass material is removed, artificial fertilizers must then be used. This is likely to be considered an unsatisfactory trade-off both environmentally and from an energy balance perspective.

Livestock residues are another special category of biomass. While there is probably the equivalent of around 20–40 EJ of livestock residue generated each year, most of this is in the form of dung, which has a very low energy content and is not a cost-effective fuel for power generation (though it is used for heating in some parts of the world). It is only where livestock is farmed intensively that it becomes economical to utilize the waste, and then only when the operation is being carried out on a sufficiently large scale.

Dairy and pig farms fall into this category and it can be cost effective to use a biomass digester to convert the animal effluent into a biogas containing methane that can be burned in a gas engine to generate power. It is often not cost effective for small farms to install their own digesters, but in Denmark there are schemes where the waste from a number of small farms is collected and then processed centrally. Sewage farms that treat human waste are another potential source of methane-rich gas. Meanwhile, poultry farm residues have been used in combustion plants in the United Kingdom and United States.

Wood waste comes from three sources: material that can beneficially be removed from natural and managed forests to improve the health of the plantation, residues left in a forest after trees have been logged, and the waste produced during the actual processing of wood in sawmills and paper manufacturing plants. Process plant waste is the cheapest and most economical to utilize. Many sawmills and most modern paper plants burn their waste, producing heat and electricity for use in the facility. Any surplus power may be sold. Residues left after logging are generally expensive to collect and transport, but they have been utilized in situations where the demand for biomass fuel is high. Similarly, the removal of dead trees and undergrowth from natural forests, while improving their health and reducing the risk of fire, is an expensive process that only becomes cost effective if the value of the fuel is high.

FUELWOOD

Fuelwood is the wood that is used for heating and cooking in many parts of the world and that still makes up a significant proportion of primary energy consumption, particularly in poorer areas of the world. Such usage is expected to continue well into the 21st century. In sub-Saharan Africa, for example, these fuels will still account for 42% of primary energy consumption in 2035 according to the International Energy Agency.

The use of fuelwood is also important in Asia, although consumption in India and China, two major users in the past, is expected to fall over the next two decades. Latin America and the Caribbean also use significant quantities. Many regions also use large quantities of charcoal, which is produced using traditional methods from fuelwood. Total fuelwood consumption was around 18 EJ in 2005 according to the World Energy Council.

In the future some of this wood could be utilized for power generation. The amount that would be available depends on the sustainability of the forests from which it is taken, but it could represent a significant resource for future renewable generation.

ENERGY CROPS

The large-scale development of biomass power generation will depend on a supply of fuel from energy crops. Wastes from the sources outlined above can only provide fuel for a limited generating capacity. Large-scale bioenergy farming is already carried out for the production of biofuels; sugarcane is produced in Brazil and maize in the United States, both for ethanol production, while oil-seed crops are grown in Europe and palm oil in the tropics, both of which provide oils. These liquid fuels are primarily used for transportation applications with ethanol added to gasoline and bio-oils used as diesel substitutes. Crops of a different sort are required for power generation. Such crops are being grown on a small scale in some parts of Europe and the United States, but if a significant industry is to be developed, then much larger areas of land will need to be devoted to such production.

An energy crop that is suitable for a power plant fuel must be fast growing, provide a high annual yield, and be relatively easy to harvest. A number of species have been tested and two groups appear to show the most promise for this purpose: fast-growing trees and grasses. Both of these can be grown relatively easily with minimal use of fertilizers, which are expensive, and both demand relatively little management once established.

Among the tree species that appear to show promise are several species that have traditionally been coppiced as a source of fuel domestic. Willows and poplars thrive in Europe, Scandinavia, and parts of the United States and Canada, and both species can be coppiced.

Coppicing is still the preferred means of harvesting these trees as power plant fuel. Once plantations are established, the growth is cut to the ground every three to five years leaving a stump from which new growth will appear. In warmer, more southerly climates other species such as sweetgum, cottonwood, sycamore, and eucalyptus have also shown promise. Some of these must be harvested using more traditional wood-cutting techniques.

Since wood from a single plantation cannot be harvested every year, a crop industry will require plantations that can be harvested in rotation so that fuel is available all the time. Typical yields from wood plantations are 10–11 dry tonnes/ha/y for willow and 8–13 dry tonnes/ha/y for poplar as shown in Table 15.1. This compares with a yield of 2.5 dry tonnes/ha/y for forest biomass.

Wood harvesting requires specialist equipment and new types of machinery are being developed for the purpose. Wood, when harvested, contains around 50% moisture and should be dried before combustion to offer the best energy

TABLE 15.1 Energy Crop Yields

	Yield (Dry tonnes/ hectare/year)
Switchgrass	7.7–14.3
Hybrid poplar	8.1–12.8
Willow	10.1–11
Forest biomass	2.5

Source: U.S. Department of Agriculture.

content. Once dried, wood can be burned directly without need for processing other than sawing. However, wood is sometimes ground to sawdust and formed into pellets using a binder. The resulting fuel is more expensive than untreated wood but does offer a standard product for combustion.

Grasses grow rapidly each year, providing an abundant supply of biomass that can be harvested annually in the autumn when the stalks are virtually dry. Under these conditions the water content of grasses is around 15%, much lower than for freshly harvested woods. Grasses that show promise as crops include switchgrass and miscanthus. The grasses will grow in a variety of habitats but are best suited to natural grasslands such as the U.S. prairies or grassy regions in other parts of the world. Annual yields for switchgrass are 8–14 dry tonnes/ha, as shown in Table 15.1.

The harvesting of grasses is more straightforward than for wood and can be carried out with harvesters and balers similar to those used for cereals. The energy density of loose grass is much lower than wood, so grass is generally formed into briquettes or pellets before transportation. In this form, grass will have around 95% of the energy density of wood.

Table 15.2 shows figures for the calorific value of a range of biomass fuels. As-harvested wood contains around 10 GJ/tonne, rising to 19 GJ/tonne once it has been dried. Cereal straw contains 15 GJ/tonne, while the energy content of switchgrass is 16 GJ/tonne and for miscanthus it is 16 GJ/tonne. Undried, coppiced willow has an energy content of 10 GJ/tonne typical of an undried wood. In comparison, bituminous coal generally contains 27–30 GJ/tonne and oil 42–45 GJ/tonne.

The combustion of grasses normally produces more ash than the combustion of wood. Against this, the content of trace elements such as chlorine, nitrogen, potassium, and sulfur are often lower, making emission control simpler.

Growing and harvesting a biomass combustion fuel is only the first stage of complex infrastructure that will be required to support the long-term and large-scale development of biomass for power generation. Since harvesting will often be seasonal, fuel must be stored to provide a year-round supply and

TABLE 15.2 Calorific Value of Biomass Fuels

Fuel	Calorific Value (GJ/tonne)
As-harvested wood	10
Dry wood	19
Straw	15
Miscanthus	19
Coppiced willow	10
Switchgrass	16
Bituminous coal	27–30
Oil	42–45

Source: U.K. Energy Technology Support Unit; *Biomass Energy Crops: Massachusetts' Potential*, Massachusetts Division of Energy Resources and Massachusetts Department of Conservation and Recreation, 2008; and Oak Ridge National Laboratory.

transportation networks need to be established to keep power plants supplied. This will require extensive integration of the agriculture and energy industries.

The key issue to be solved, if development is to proceed, is how much land can be put aside for fuel production and how much must be retained for food production. Difficulties are already arising in the tropics where rain forests are being destroyed to make room for palm oil plantations. Balancing the two will be difficult, but there does appear to be areas of land not used or not suitable for food production that could provide combustion fuel. These include some northern forest regions.

The areas of land required to supply a biomass power plant are considerable. On the basis that one hectare of land can produce 10–12 tonnes of dry fuel each year, a 10 MW power plant would require around 7000 ha dedicated to supply it with fuel.

BIOMASS TRADE

The relatively low energy density of biomass as combustion fuel means that it is generally not considered economical to move it over long distances. In spite of this, there is a growing trade in biomass combustion fuel in the form of pellets, particularly in Europe. The driving force for the trade is the push toward renewable generation in Europe and the various subsidies available to support it. This includes support for biomass power generation in the form of feed-in-tariffs of tax relief that can make it cost effective to purchase wood pellet fuel for a biomass power plant.

The main sources for wood pellets in Europe are the United States and Canada with a growing supply from Russia. There is a nascent export industry in Australia and New Zealand, and South Africa also supplies pellet fuel to Europe. Meanwhile, Canada sends pellets to Japan, a country with few natural energy resources that must import all of its fossil fuel as well.

The size of the pellet industry remains small but it has the potential to grow if current incentives continue. There is also potential for a market similar to that in Europe opening in the United States. There are two primary sectors of the pellet market: the residential and the industrial. In the United States most pellets are consumed residentially, but in Europe the big market driver is industrial use by power utilities.

BIOMASS ENERGY CONVERSION TECHNOLOGY

There are a number of ways of converting biomass into electricity but the most common is to burn the material in a furnace, raising steam that is used to drive a steam turbine, an approach analogous to the use of coal in a coal-fired power plant. The main alternative to this is biomass gasification in which the fuel is converted into a combustible gas that can be burned to provide heat. The development of coal gasification systems is expected to drive improvements in biomass gasification.

Biomass plants tend to be small compared to conventional coal-fired plants and less efficient. However, efficiency can be improved if they can provide heat as well as electricity in a combined heat and power plant, and this configuration is common, particularly in industries that use their own biomass waste for power generation. A further method that allows for more efficient use of biomass is co-firing. This involves adding a proportion of biomass to the coal in a coal-fired power plant. Large modern coal plants can operate at high efficiency and, when co-fired, biomass is converted into electricity with similar efficiency.

For some animal wastes it is also possible to generate a combustible gas using anaerobic digestion. This is the same process that occurs in landfill sites and generates methane. Digesters can be designed to generate methane from both animal and human wastes. The gas is then normally used to fire a gas engine to provide electric power.

DIRECT FIRING

The direct firing of biomass involves burning the fuel in an excess of air inside a furnace to generate heat (Figure 15.1). Aside from heat the primary products of the combustion reaction are carbon dioxide and a small quantity of ash. The heat is absorbed by a boiler placed above the main furnace chamber. Water flows through tubes within the boiler where it is heated and eventually boils, producing steam that is used to drive a steam turbine. Direct-firing technology was developed in the 19th century for coal combustion but has been adapted to other

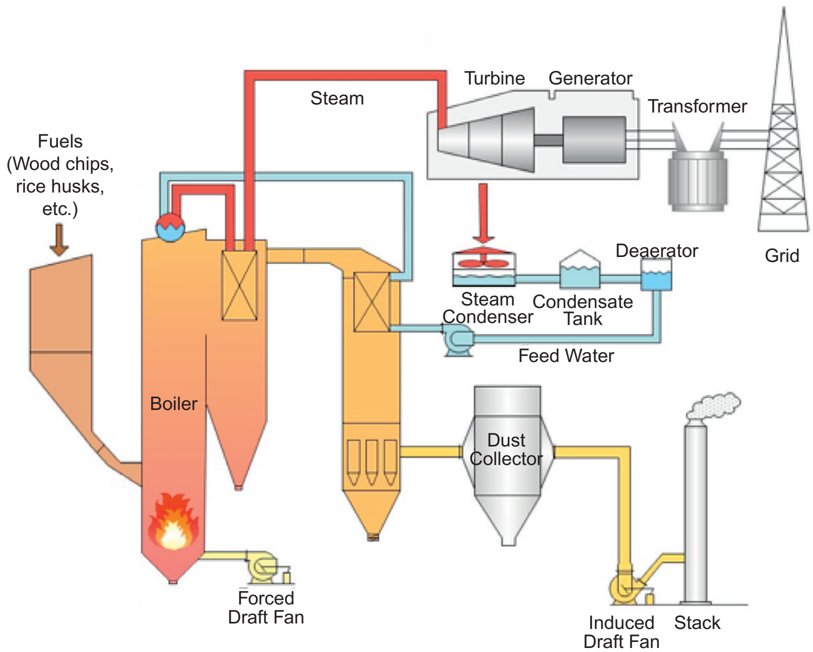

FIGURE 15.1 Schematic of a direct-fired biomass power plant.

fuels including biomass. While the heat from a biomass furnace is normally used to generate steam in this way, it may be exploited directly in some industrial processes too.

The simplest type of direct-firing system has a fixed grate onto which the fuel is piled and burned in air that enters the furnace chamber from beneath the grate (underfire air). Further air (overfire air) is then added above the grate to complete the combustion process. This type of direct-firing system, called a pile burner, can burn wet and dirty fuel but its overall efficiency is only around 20% at best. The fixed grate makes it impossible to remove ash except when the furnace is shut down, so the plant cannot be operated continuously either, a further disadvantage for this design.

An improvement over the pile burner is the stoker combustor. This type of combustor allows fuel to be added continuously, either from above (overfeed) or from below (underfeed), and has a mechanism for continuously removing ash. The stoker grate was first developed for coal combustion in the 1920s and in the 1940s the Detroit Stoker Co. designed a stoker boiler for wood combustion.[3]

3. *Biomass Combined Heat and Power Catalog of Technologies*, U.S. Environmental Protection Agency Combined Heat and Power Partnership, 2007.

Fuel is distributed more evenly in a stoke grate than in a pile burner, allowing more efficient combustion. Air still enters the furnace from beneath the grate and this air flow cools the grate. The airflow determines that maximum temperature at which the grate and thus the furnace can operate and this in turn determines the maximum moisture content of the wood that can be burned, since the dampest wood will require the highest temperature if spontaneous combustion is to be maintained.

For an underfeed stoker combustor, fuel is added from beneath and the grate behaves like a slowly erupting volcano. Ash, when it is formed, falls down the flanks of the fuel pile and is removed from the sides. In an overfeed stoker combustor, the grate itself normally moves via some form of chain mechanism allowing fuel to be added from one side and ash removed from the other (Figure 15.2). This type of design is called a mass-feed stoker. A second type, called a spreader stoker, disperses finely divided fuel pneumatically across the whole surface of the grate with finer particles burning in the region above the grate. There are a number of other refinements to the stoker combustor such as an inclined and water-cooled grate. Even so, maximum overall efficiency is only 25%.

Most modern coal-fired power plants burn finely ground coal that is fed into the power plant furnace through a burner and then ignites in midair inside the furnace chamber, a process called suspended combustion. It is possible to burn biomass in this way but particle size must be carefully controlled and moisture content of the fuel should be below 15% (Figure 15.1). Suspended combustion of biomass, while it can provide a higher efficiency, is not widely used in dedicated biomass power plants. However, it does form the basis for co-firing, which is discussed at greater length in the following section.

The main alternative to the stoker combustor for modern direct-fired biomass plants is the fluidized bed combustor. This type of combustor can cope

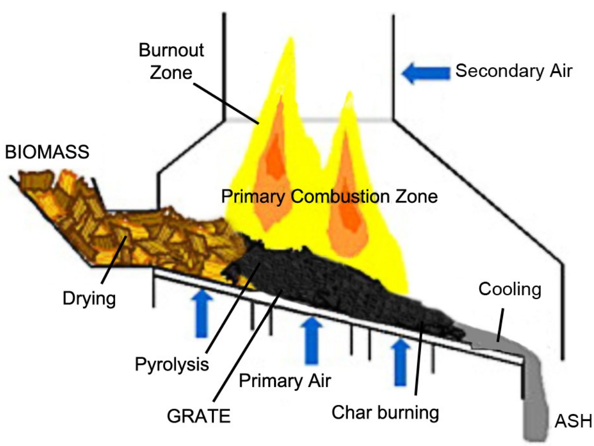

FIGURE 15.2 Cross-section of a moving grate for biomass combustion.

with fuels of widely differing type and quality, making it much more versatile than a stoker combustor. The fluidized bed contains a layer of a finely sized refractory material, such as sand, that is agitated by passing air through it under pressure so that it becomes entrained and behaves much like a fluid. Fuel is mixed with this refractory material where it burns (when the bed is at its operating temperature) to release heat as in a conventional furnace. Depending on the pressure of the air that is blown through the fluidized bed it will be a bubbling bed, behaving much like a boiling fluid, or a circulating bed in which the particles are entrained with the air and those that escape the boiler are subsequently captured in a cyclone filter and recycled. Fuel content within the bed in usually maintained at around 5%.

Fluidized beds can burn a wide range of biomass fuels with moisture content as high as 60%. For low-quality fuels the bubbling bed is preferred, whereas the circulating bed is better for high-quality fuels. Overall efficiency is again only 25% at best, similar to a stoker combustor.

Where the fuel quality is low, gas or coal can be used to raise the bed temperature sufficiently at startup for combustion to commence. An additional advantage of the fluidized bed is that material can be added to the bed to capture pollutants like sulfur that would otherwise result in atmospheric emissions.

Direct-fired biomass power plants typically have a generating capacity around 25–50 MW. This small size, combined with the relatively low combustion temperature in the furnace (biomass is more reactive than coal and so tends to burn at a lower temperature), are the two main reasons for these plants' low efficiencies compared to coal plants where overall efficiencies above 40% are now common in new facilities.

Improvements are possible. Increasing the size of the typical plant to 100–300 MW would allow larger, more efficient steam turbines to be used, and several 300 MW are being planned in Europe. New small steam turbines that incorporate advanced design features currently found only in large coal plant turbines will also improve efficiency. Adding the ability to dry the biomass fuel prior to combustion can result in a significant increase in performance. With these changes, direct-fired biomass plants should be able to achieve 34% efficiency.

CO-FIRING

Much more efficient conversion of biomass into electricity can be achieved quite simply and on a relatively large scale in another way—by the use of co-firing. Co-firing involves burning a proportion of biomass in place of some of the coal in a coal-fired power plant. Since most coal stations operate at much higher efficiencies than traditional direct-fired biomass plants, co-firing can take advantage of this to achieve conversion efficiency of 40% or more in a modern high-performance coal-fired facility.

There is another form of co-firing in which a predominantly biomass combustion plant uses a fossil fuel, normally natural gas, to both stabilize and supplement the biomass fuel. This technique will normally be used in a dedicated biomass plant to increase performance and flexibility.

Co-firing of the first type is attractive to coal plant operators because it allows them to burn biomass and therefore reduce their net carbon dioxide emissions with very little plant modification. Biomass co-firing can also reduce sulfur emissions because biomass contains virtually no sulfur. Since coal-fired plants can burn large quantities of biomass this also offers a means of establishing the biomass infrastructure needed to enable biomass power generation to develop into a large-scale industry.

The most efficient and common type of coal-fired power plant in operation is the pulverized coal (PC) plant that burns coal that has been ground to a fine powder. Plants of this design can burn up to 10% biomass with little modification to their plant. Biomass is simply mixed with the coal before it is delivered to the coal mills where the mixture is ground prior to injection into the combustion chamber. For a 1000 MW power plant this would be equivalent to a 100 MW biomass plant, but with much higher efficiency than a dedicated plant.

Simple co-firing of this type is limited by the type of biomass fuel that can be used and by the proportion of co-firing possible. To overcome both these limits another approach is to have a dedicated biomass fuel delivery line. Fuel from this line can either be mixed with the powdered coal before combustion or delivered to dedicated biomass burners in the furnace. The latter are usually located lower down the combustion chamber allowing a longer transit time to completely burn the biomass fuel. This will allow a plant to burn up to 15% biomass by heat content. Higher proportions are possible but generally require greater plant adaptation and consequently more expense.

Co-firing has become common in the United Kingdom where most plants burn 5–10% biomass and is becoming so in the United States with plants generally using 10% biomass. At this level, boiler efficiency can fall by up to 2%. Such plants will typically burn wood wastes and some will use wood pellets.

An alternative to conventional co-firing involves gasification of the biomass in a dedicated biomass gasifier attached to the coal plant. The combustible gas is then burned in the coal-fired furnace. This approach is more expensive than the simple co-mixing of fuels but it avoids some of the problems that can be associated with conventional co-firing and allows a greater proportion of biomass to be burned.

A third approach is called parallel co-firing. This involves having a separate biomass furnace, with the hot gases generated by combustion of biomass being mixed with the hot coal-flue gases. In a variation of this approach, the biomass has its own steam-raising boiler too with the steam flows blended before entering the steam turbine. Again this is more costly than conventional co-firing.

Conventional co-firing remains the favored approach but it does have its problems. Some wastes such as sawdust can block fuel feed systems and need to be avoided. Other biomass fuels such as grasses have a high alkali content that can cause problems in coal-fired boilers. Biomass is also more volatile than coal when burned so much of the combustion takes place higher in the combustion chamber and more overfire air may be necessary to ensure complete combustion. Additionally, the ash from a co-fired boiler has a different composition to that from a plant that burns only coal. Coal plant ash is often used in various ways by the building industry but the reuse of the ash from a plant that burns biomass has led to regulatory difficulties in the past.

BIOMASS GASIFICATION

Gasification involves the heating of biomass in a reducing atmosphere with the addition of water vapor to generate a gas that contains around 40% hydrogen and carbon monoxide with most of the remainder being nitrogen. The composition of a typical wood gas from a gasifier is shown in Table 15.3. Further reaction can convert the carbon monoxide into more hydrogen if required and the gas will usually be cleaned before combustion. Wood gas has a low calorific value between 4 MJ/m^3 and 19 MJ/m^3; by comparison the calorific value of natural gas is around 38 MJ/m^3.

There are two types of gasifier used for biomass gasification: a fixed bed and a fluidized bed. The fixed bed is the simplest type, comprising a cylindrical vessel with a grate at the bottom, as shown in Figure 15.3. Biomass is introduced at the top and air from the bottom, and the heat for the gasification process is provided by the partial combustion of the biomass. Ash is then removed from the bottom of the gasifier. Fixed-bed gasifiers are relatively primitive and inefficient and are generally only used at capacities of 5 MW of less.

The alternative, a fluidized-bed gasifier, is similar to the fluidized-bed combustor described previously but is operated with a reduced supply of air so that combustion cannot go to completion (Figure 15.4). The heat content of the gas from both types of gasifier is often less than 6 MJ/m^3. The low-quality gas from

TABLE 15.3 Composition of Wood Gas

Nitrogen	51%
Carbon monoxide	27%
Hydrogen	14%
Carbon dioxide	5%
Methane	3%
Oxygen	1%

Source: Penn State Energy Institute.

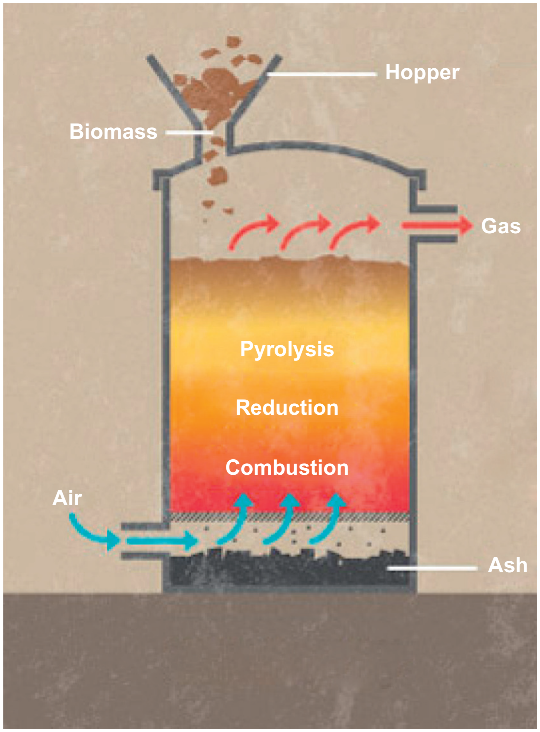

FIGURE 15.3 Cross-section of a fixed-bed biomass gasifier.

a fixed-bed gasifier has limited use and will normally be burned to raise steam to drive a steam turbine. The gas from a fluidized-bed gasifier can potentially be utilized in a gas turbine, a reciprocating engine, or a fuel cell too.

One of the potentially most interesting gasification configurations is the biomass-based integrated gasification combined cycle (IGCC) plant. Plants of this type have been developed for coal combustion and form one of the future options for carbon dioxide capture and storage from coal-fired power stations. In a biomass IGCC plant, biomass is gasified, cleaned, and burned in a gas turbine. The hot exhaust gases from the gas turbine are then used to raise steam to drive a supplementary steam turbine. With close integration of the plant components, an overall efficiency of 45% might be achieved. A further advantage of this type of plant is that since biomass is carbon neutral, no carbon dioxide capture is necessary.

Alternative ways of exploiting low-quality biogas include the use of an organic Rankine cycle turbine similar to that used to generate electricity from low-grade geothermal reservoirs or of a Stirling engine similar to that used in some solar thermal plants. If the mixture of carbon monoxide and hydrogen is converted into just hydrogen it could also be used in a fuel cell. Some

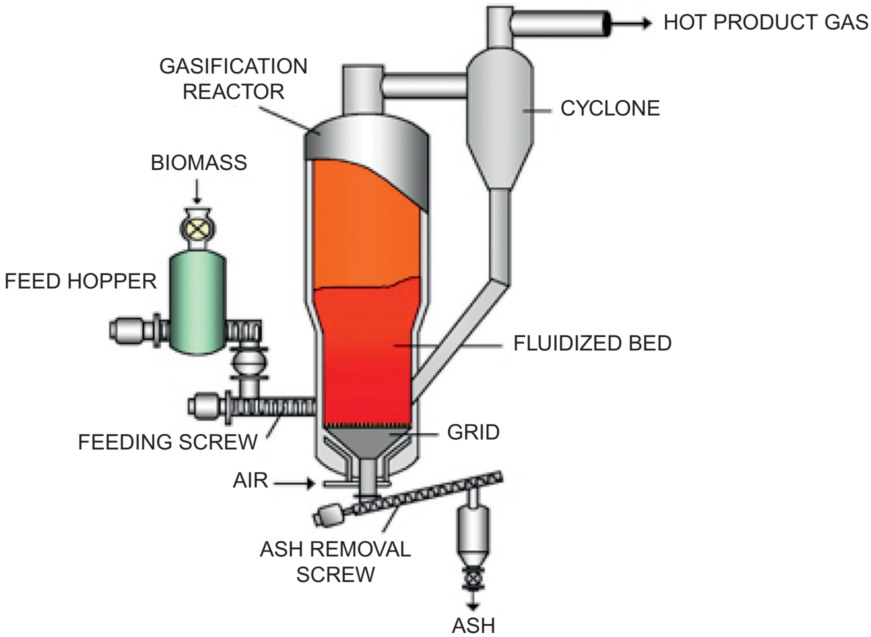

FIGURE 15.4 Schematic of a fluidized-bed biomass gasifier.

high-temperature fuel cells can burn carbon monoxide as well as hydrogen, making this conversion potentially unnecessary.

As outlined earlier, another way of exploiting biomass gasification is to feed the output of the gasifier directly into the combustion chamber of a coal-fired power plant, a form of co-firing. This avoids many of the problems associated with utilizing a low-grade gas but is an expensive co-firing configuration.

FUEL HANDLING

Fuel handling is important for an effective biomass power plant. Since most biomass fuels are seasonal, there must be facilities to store large quantities if power plants are to be supplied throughout the year. This is particularly significant for grasses that will all be harvested in the autumn. These will normally be formed into briquettes and then stored until needed, but if this is to be the primary source of fuel for a power plant, a store capable of holding one year's supply will be needed.

Woods can be harvested at different times of the year and this can help with the fuel supply management. Depending on the harvesting techniques the wood will arrive at the power station or storage depot in the form of bundles of coppiced branches, chipped wood, or whole trees. Wood is easier to chip when green so this will normally be carried early in the harvesting process, probably before it reaches the power plant. The latter will generally need to be able to store several weeks' supply of fuel under conditions where it is protected from the weather so that it does not get wet. Otherwise, the fuel will deteriorate and lose energy content.

The efficiency of a biomass power plant will depend on the moisture content of the fuel. The lower the moisture content, the higher the efficiency. For example, a reduction in moisture content from 50%, as harvested, to 10% after storage can result in boiler efficiency rising from 70% to 83%. It is important, therefore, to allow wood to dry before it is burned.

The form in which fuel is actually supplied to the power station will vary. Most often it will be chipped although the use of pellets is expanding. For plants that use suspension firing (e.g., in co-firing with coal) the chips must be further reduced in size using grinding equipment. There is also a novel approach in which whole trees are supplied to a specially designed power plant. The trees are dried for 30 days and then sections are delivered to the boiler that is designed especially for this form of fuel. Efficiencies of 34% have been predicted for a 150 MW plant of this type.[4]

BIOMASS DIGESTERS

Biomass will ferment naturally in the absence of air to produce a gas that is rich in methane. This is the process that occurs underground in landfill waste disposal sites and it can also be found in the lakes of large hydropower plants when there is a large quantity of biomass immersed beneath the waters. The same process can be harnessed in biomass digesters to produce a methane-rich gas from wastes. The technique is normally applied to agricultural animal wastes.

Biomass digestion is only cost effective for large farming operations, most usually on dairy or pig farms where the slurry produced by the animals must be treated to prevent it causing an environmental hazard. Digesters of differing sophistication are available depending on the size of the farm. For small farms the most suitable is usually a lagoon digester, essentially a pond (the lagoon) into which the slurry is placed. The lagoon is covered with an impermeable membrane cover that is used to collect the emitted gas. The slurry must contain less than 2% solids and the lagoon must usually be maintained above 30°C, which limits the application to warmer climates since it is not economic to heat a lagoon digester.

A more sophisticated system is the tank digester (Figure 15.5). Slurry is loaded into a tank that is fitted with a stirring mechanism to mix the contents evenly. The tank can be heated to keep the fermentation at the optimum temperature. Tank digesters can handle slurries with 3–10% solids. For slurries with higher solids content a plug flow digester is preferable. This has three elements:

4. K. W. Ragland, L. D. Ostlie, D. A. Berg, *Whole Tree Energy Power Plant*, 2005.

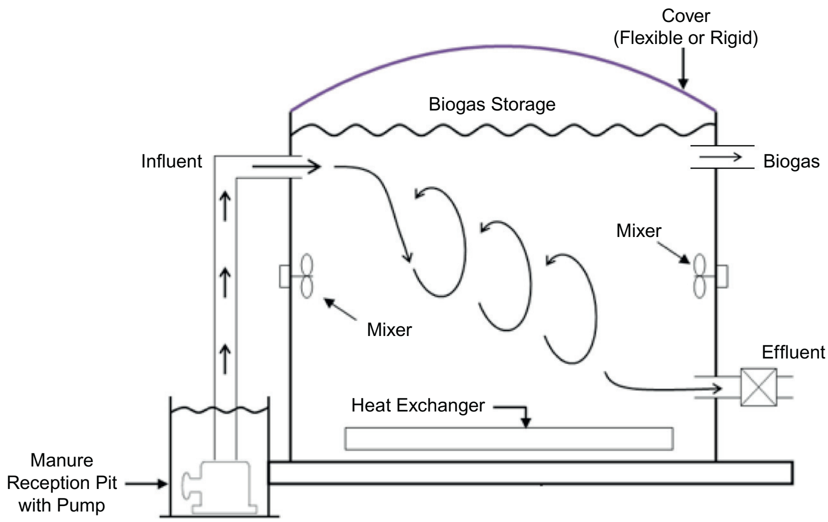

FIGURE 15.5 Biomass tank digester.

a mixing tank, digester tank, and settling tank. The slurry is first fed into the mixing tank from where it enters the digester tank, which contains heating pipes to maintain the ideal temperature. The material moves slowly across the digester tank, with fermentation proceeding to completion in about 20 days as it crosses the tank. After 20 days it passes into a settling tank where the remaining solid material is removed and can be used as fertilizer.

The gas from an anaerobic digester has a heating content of 22 MJ/m^3, suitable to be burned in a reciprocating engine to generate electricity and heat. However, the capital cost of such systems is high and can only be supported when there is a large quantity of waste to ferment. Similar systems can be used to treat municipal sewage waste and they form an effective means of both rendering it harmless and producing a valuable by-product.

Most biomass digester-based power generation plants are relatively small with capacities of 100 kW or less. Landfill gas sites, which can produce large volumes of methane, can sometimes support gas engines with generating capacities of 20–30 MW.

LIQUID FUELS

There are a range of liquid biofuels that are manufactured from biomass sources. These include bio-alcohols made from the fermentation of crops rich in starch or sugar and biodiesel, which is derived from oil-producing crops.

One of the most important liquid biofuels today is ethanol, which produced in Brazil from sugarcane and in the United States from maize. In both cases the fuel is blended with gasoline so that up to 10% of vehicle fuel may be bioderived. There is a growing ethanol production industry in Europe too where other crops such as wheat may be used as the feedstuff. Some biomethanol is also produced.

Europe is the main region in which biodiesel is produced and used. The fuel is derived from crops such as sunflowers or oil-seed rape. It can also be produced from animal wastes and some palm oil is imported into Europe for this use. As with ethanol, biodiesel is blended with diesel fuel for transportation. A European Biofuel directive from 2003 called for 5% of biofuel to be blended with diesel by 2010. By 2012 the actual proportion was 4.5%. A further directive intended that this should rise to 10% by 2020, but a recent policy change means that food-derived biofuel, which accounts for most of that produced, will be limited to 5%. This change reflects the still unresolved issue of how to balance fuel and food production.

Ethanol and biodiesel are considered first-generation biofuels. Scientists are currently trying to develop second-generation fuels that are made from nonfood sources such as cellulose and algae. If these processes can be developed effectively, then the impact of liquid biofuel production on food production should be much smaller.

While most of the liquid biofuel manufactured is being used for transportation fuel, both types of biofuel can also be used to generate electric power. The most suitable generators are reciprocating engines, but gas turbines can also burn biofuels and they can be used to fire a boiler to raise steam.

COST OF BIOMASS POWER GENERATION

A biomass-fired power station is technically similar to a coal-fired power plant and the economics of the two are based on similar principles. In both cases the cost of the electricity generated depends on two factors: the cost of building the plant and the cost of operating the plant. The first of these is usually dominated by the cost of the actual installation, although the cost of any loan required to finance the project can also have a significant impact. The second depends mostly on the cost of the fuel.

Technology Costs

The capital cost of biomass installations varies widely. The cheapest option for generating electricity from biomass is co-firing. Retrofitting a co-firing option to an existing coal-fired power plant costs between \$50/kW and \$500/kW of biomass generating capacity in the United States depending on the type of boiler. Prices are likely to be similar in Europe and these are the two regions where co-firing is most popular.

For direct-fired biomass systems, the cost depends on the size of the plant and whether it is for power generation alone or for combined heat and power generation. Figures from the U.S. Environmental Protection Agency (EPA)[5] from the last decade put the cost of a 500 kW stoker boiler for CHP use at \$9300/kW. For an 8 MW unit this falls to \$4000/kW, while the same unit used for power generation alone has a cost of \$1600/kW. Overall combustion power plant costs have risen steeply since these estimates were made and the U.S. Energy Information Administration (EIA) estimate for the capital cost of direct-fired biomass plant for power generation alone of \$1700/kW in 2007 had risen to \$3400/kW by 2011. Meanwhile, the U.S. EPA analysis suggested that biomass gasification was roughly twice as expensive as a direct-fired plant. Anaerobic digesters are also much more expensive than direct combustion plants, but plants designed to generate from landfill gas or waste treatment plants can be slightly cheaper than the direct-fired combustion plant.

Fuel Costs

The cost of biomass fuel depends on its source. Some agricultural and industrial wastes cost nothing. The same applies to landfill gas. In most cases, however, power plant operators will have to pay for their fuel whatever its source. Indeed, many of the best wastes are being converted into biomass fuel, such as pellets, that are then sold at a premium.

From a long-term perspective fuel crops are likely to be the most important source of biomass fuel. In Europe, particularly in the United Kingdom, wood chips are being used as biomass power plant fuel. Typical cost is around £100/tonne at 35% moisture content. In comparison, the cost of wood pellets is roughly £200/tonne (but energy content of the wood chips with 8% moisture content is significantly higher than wood at 35% moisture content). The fact that the latter is becoming a commodity means that prices are subject to greater competition and therefore could fall in line with other biomass fuel prices.

The other potential future energy is grass. Although little grass for biomass fuel is grown, estimates indicate that it may have a slightly lower cost per tonne to wood chips. However, the slightly lower energy content means that costs are likely to be more or less the same.

Cost of Electricity from Biomass

Electricity generation costs from biomass plants depend on the type of plant. Some, though of limited application, can have very low costs. For example, in California the electricity produced by the anaerobic digestion of food waste can be generated for about 10% of the cost of power from a direct-fired biomass plant. Power generation from landfill gas methane and waste treatment plants is also cost effective where the energy source is available.

5. *Biomass Combined Heat and Power Catalog of Technologies*, U.S. Environmental Protection Agency Combined Heat and Power Partnership, 2007.

For the more conventional, and more widely applicable, direct-firing biomass power plant, the U.S. EIA has estimated that the levelized cost of electricity in 2011$/MWh for a plant entering service in 2018 is $110/MWh.[6] Meanwhile, estimates by the Oregon Department of Energy put the cost of generation today between $52/MWh and $67/MWh in the Pacific northwest. In the United Kingdom electricity costs of £35–40/MWh have been proposed by some wood chip suppliers.

In all cases the actual cost of generation will also depend on the level of subsidies available for electricity generated from renewable sources.

6. Levelized Cost of New Generation Resources in the Annual Energy Outlook 2013, U.S. Energy Information Administration, 2013.

Chapter 16

Power from Waste

The generation of power from waste is a very specialized industry. Power-from-waste plants are designed to burn residual urban waste that cannot be recycled to reduce its volume, destroy potentially hazardous materials, and generate heat and power. The plants include extensive environmental controls to ensure that they do not release any toxic emissions and are consequently expensive to build. However, their economics do not depend on the value of the electricity they produce since the operators are paid for the volume of waste they process. In this sense electricity and heat are valuable by-products.

Urban waste is a major problem throughout both the developed and the developing world but it is in the advanced economies of North America, Europe, and some nations in southeast Asia that the problem is greatest. These economies generate large volumes of glass, metal, paper, plastic, and organic waste that must be collected, sorted, and then treated or stored in some way. At the top of the national league is Canada, which produces around 780 kg for each person every year. The average amount of waste generated per person per year in the European Union (EU) is 577 kg, and some individual EU nations such as Ireland are at the upper end of national rates of waste production. At the other end of the scale of advanced countries is Japan with a per-capita output of less than 400 kg. In contrast, a typical town dweller in Sub-Saharan Africa may produce less than 100 kg each year.

Once urban or municipal waste has been collected, there are three primary ways in which it can be handled, as shown in Figure 16.1. The first, and perhaps the most important today, is to sort it and then recycle everything that can be reused. This should include the composting of organic waste so that it can be returned to the soil. The second option is to burn the waste. Not all will be combustible but a large part of urban waste can be burned including parts that should ideally be recycled. The final option is to bury it in a landfill site. This is the easy option and was preferred in the past in many countries, but pressure on land coupled with greater environmental awareness is leading to a reduction in the use of landfills in most advanced economies.

The level of exploitation of waste-to-energy plants varies from country to country. They have been used widely in parts of Europe, where waste has been burned since the end of the 19th century, and form a major part of Japan's waste disposal strategy. In contrast, the United Kingdom and United States have only adopted the technology patchily, while some nations, such as Ireland and Greece,

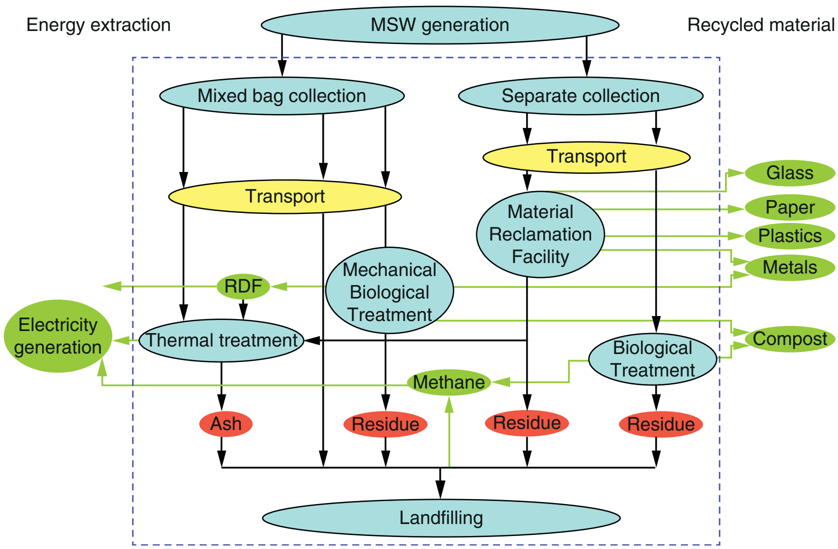

FIGURE 16.1 Urban waste collection and disposal.

do not incinerate any waste. This is partly tradition, but environmental concerns about the emissions from the plants have resulted in local resistance to their construction in some parts of the world. More advanced emission treatment processes may make it easier to build these plants in the future.

Some countries have used the incineration of waste extensively. In Sweden, for example, 14% of municipal waste is put into landfill sites, 41% recycled, and 45% is incinerated. By contrast, the United Kingdom disposes of 74% of its waste in landfill sites, recycles 18%, and incinerates 8%. The Netherlands, meanwhile, incinerates 33%, recycles 64%, and only disposes of 3% in landfill sites. Within the EU, Denmark has the highest incineration rate of 54%.

Where they are employed, power-from-waste plants generally burn domestic and urban refuse—called in this context municipal solid waste (MSW)—using the resulting heat to generate steam to drive a conventional steam turbine. MSW can also be sorted and treated to produce a compacted fuel called refuse-derived fuel (RDF) that can be burned in a power station. Some industrial waste may be treated in the same way. However, industrial wastes are likely to contain toxic materials that have to be handled using special procedures. Where such care is not required, they can be dealt with in the same way as urban waste.

In addition to municipal waste and industrial waste there are a number of other categories of waste, primarily resulting from the agricultural and forestry industries, that can be used to generate electricity. These have been dealt with under biomass in Chapter 15. That chapter also dealt with the collection and use of methane produced in landfill refuse disposal sites. However, we need to consider landfill briefly here since it offers the main alternative to waste combustion for all urban waste that cannot be recycled.

LANDFILL WASTE DISPOSAL

The landfill site—essentially an enormous hole in the ground (or more accurately a natural depression since sites are not generally excavated first)—is the main alternative to the technologies discussed in this chapter as a means of waste disposal. Though crude, its simplicity has led it to become the favored method of urban waste disposal across the globe. Waste that has been collected is simply off-loaded at the site until the depression or hollow is considered full. At that point earth is bulldozed over the deposited waste and the whole structure is left to settle. Over time the organic material within the site will decompose, producing carbon dioxide if there is air present but methane if the decomposition takes place in the absence of air. Methane is a common product in many landfill sites and its production can continue for one to three decades after the site has been closed and sealed.

While landfill use remains popular in many countries, it is coming under pressure in others. This is partly a result of the demand for land that increasingly restricts that available for waste burial. More potent still are environmental concerns about the lack of recycling and the long-term effects of landfill disposal—effects resulting from the methane emissions from such sites and from the seepage of toxic residues into water supplies.

Such concerns have already led the EU to legislate[1] to restrict the use of landfill waste disposal. Where it does not already exist, similar legislation can be expected in other parts of the world. But waste will still be produced. This is where technological solutions such as the power-from-waste plant enter the equation.

Power-from-waste technology is not cheap. The specialized handling that waste requires, coupled with the need for extensive emission control systems to prevent atmospheric pollution, make such plants much more expensive to build than any other type of combustion power plant. They are also expensive to operate.

If these plants had to survive on the revenue from power generation alone, they would never be built. Fortunately, they have another source of income. Since waste has to be disposed of in a regulated manner, waste disposal plant operators can charge a fee—normally called the tipping fee—to take the waste. The tipping fee represents the main source of income for a power-from-waste plant. Any additional income derived from power generation will benefit the economics but the plant may well be able to survive without it.

SOURCES OF WASTE

There are two principles types of waste suitable for disposal in a power-from-waste plant: urban (primarily domestic) refuse, normally referred to as MSW, and industrial waste. Some industrial waste is broadly similar in content to

1. European Union, Council Directive 1999/31/EC of April 26, 1999 on the landfill of waste, *Official Journal of the European Communities*, pp. L182/1–19 (July 1999).

MSW and this can be treated in the same way as the latter. Other industrial waste must be dealt with differently because of the hazardous or valuable materials it contains. This chapter is only concerned with MSW and it will not deal with industrial waste except where it can be burned with MSW.

The main source of MSW is an urban community. In the developed world the waste from rural communities may be handled in a similar way but this is rare in the other parts of the world. The quantity and size of such communities is growing rapidly. In the last two generations the number of people living in cities has increased by 250–500%[2] and the trend is set to continue for perhaps another generation. In the United Kingdom, for example, 79% of the population already lived in cities by 1950 but this is expected to rise to 92% by 2030. In China only 13% of the population lived in cities in 1950 but by 2030 the proportion should reach 60%. Urban dwelling has grown particularly rapidly in South America and the Caribbean where, by 2025, 80% of the populations will be living in towns. But these regions are not unique. Urban communities are growing virtually everywhere. In 2008, for the first time, more than half the world's population, or around 3.3 billion, lived in cities and towns. By 2030 the number of city dwellers is expected to reach 5 billion. These towns and cities constitute the source of MSW.

The amount of waste these populations produces varies from country to country and from continent to continent. In general, the city dweller in an industrialized country produces far more waste than one in a developing country. Thus, a typical Californian might produce 1.3–1.4 kg each day while a city dweller in Mexico City produces only half that. A Nigerian town dweller probably produces less than 200 g of waste each day.

According to the World Watch Institute, global municipal solid waste generation at the end of the first decade of the 21st century was around 1.3 billion tonnes. By 2025 the global annual production could double to 2.6 billion tonnes. The greatest producers were the 34 members of the Organization for Economic Cooperation and Development, which produced on average 1.6 million tonnes each day. Sub-Saharan Africa, by contrast, only produces 200,000 tonnes a day.

WASTE COMPOSITION

The composition of the waste varies from place to place. In general, the waste from the urban household in an industrialized country will contain 30–40% paper and cardboard and up to 10% plastic. The proportions of these in the waste from a household in, for example, the Dominican Republic will be much lower and the Dominican household's waste will probably contain 80% food waste, whereas the proportion of food in the waste from a U.S. household may only

---

2. Mining the Urban Waste Stream for Energy: Options, Technological Limitations, and Lessons from the Field, United States Agency for International Development, 1996 (Biomass Energy Systems and Technology Project DHR-5737-A-00-9058-00).

TABLE 16.1 Energy Content of Urban Wastes from Different Regions*

Region	Energy Content (kJ/kg)
United States	10,500
Western Europe	7500
Taiwan	7500
Midsize Indian cities	3300–4600
Sub-bituminous coal	10,700–14,900

**Mining the Urban Waste Stream for Energy: Options, Technological Limitations, and Lessons from the Field, United States Agency for International Development, 1996 (Biomass Energy Systems and Technology Project DHR-5737-A-00-9058-00).*

Source: United States Agency for International Development.

be 26%.[3] More generally, organic waste accounts for more than 60% of waste in low-income countries compared to 25% of the waste stream in richer counties.

There are other important differences. The waste from households in developing countries contains a high proportion of moisture, often as high as 50%, making it difficult to burn without first reducing the moisture content by drying. In contrast, the high proportions of paper and plastic in the waste from a household in the industrial world make it much easier to burn.

All these factors affect the energy content of waste, and energy content is a crucial factor in determining the viability of a power-from-waste plant. Unless the plant can produce enough excess heat from waste combustion to raise steam, then it cannot expect to generate any electricity.

Table 16.1 provides some figures for MSW energy content from different parts of the world. U.S. waste has the highest energy content of those listed. At 10,500 kJ/kg the value is approaching that of sub-bituminous coal. European cities and prosperous Asian cities such as Taipei generate waste with around 7500 kJ/kg. Meanwhile, the waste from typical midsize Indian cities contains roughly half this amount of energy. Some factors affecting the differing energy content are regional, and others are simply a matter of affluence.

In the case of Indian cities, for example, the low energy content may not be due entirely to the quality of waste. In cities in India (but not them alone) much of the urban waste is collected by city sweepers. Such waste is contaminated with considerable quantities of stone, earth, and sand. In Bombay the amount of noncombustible material of this type in waste may reach 30%. Not only does this reduce the energy content of the waste, it could also damage a combustion system, so the design of a waste handling and disposal plant has to take its presence into account.

3. Mining the Urban Waste Stream for Energy: Options, Technological Limitations, and Lessons from the Field, United States Agency for International Development, 1996 (Biomass Energy Systems and Technology Project DHR-5737-A-00-9058-00).

Given such local variations in waste content it is vitally important, before a power-from-waste plant is built, that the waste available be carefully assessed. For that, local waste collection procedures and organizations have to be examined. The issue will be particularly important for a private sector power-from-waste plant; it is less critical if the project is being built by a local municipality.

WASTE COLLECTION AND RECYCLING

Urban refuse collection is organized in different ways in different parts of the world. In some countries it is run by municipalities, in others it is provided by private operators. Where a municipality-run waste collection is a service, the same city might build and operate its own power-from-waste plant. Under these circumstances the composition of the waste can be readily assessed and controlled if necessary.

Often, however, waste collection is carried out by private companies. The waste that these companies provides will vary in quality. In some cases it will contain the whole range of waste, but in others it will have been sorted to remove the more valuable material. Some countries now require that glass, metal, plastic, and paper be recycled. This too will affect the quality of the MSW available.

Inevitably the quality of waste will vary by season. Economic factors are also important. Waste will be poorer in a recession than in a boom. Local variations can also be significant. Richer neighborhoods tend to produce better-quality waste than poorer neighborhoods. This has led to the suggestion that the quality of waste for a power-from-waste plant might be maintained by collecting only from prosperous areas of a city.

Whatever the strategy, knowledge of the waste, its source, and its variations will form a necessary part of the management of a waste-to-energy plant. That information can only be gained with practical experience, by analysis of waste collected by the contractor that will provide waste for the plant. Even with this knowledge, it may be impossible to maintain an adequate energy content in the waste throughout the year. Then the only solution may be to add some higher energy–content fuel to the waste. Biomass waste from local sources will often be the most economical solution in this situation.

Recycling is becoming an important part of municipal waste handling. This is partly a result of legislation such as that in the European Union and partly a matter of economics. World Bank figures have suggested that the global market for scrap metal and waste paper from municipal waste is worth \$30 billion each year and the total waste management market may be valued at \$400 billion annually. Even so, there is a long way to go before global waste management achieves adequate levels of recycling and reuse of material. However, the sorting of waste for recycling means that the best combustible waste that remains can also be separated, which can be valuable for a power-from-waste plant. Sorting can also help with production of RDF.

WASTE POWER GENERATION TECHNOLOGIES

A power-from-waste plant (also commonly known as a waste-to-energy or WTE plant) is a power station fueled with urban waste. As already indicated, such a facility may have as its primary function waste disposal. Nevertheless, the technologies employed will be traditional power generation technologies as used in combustion plants. Combustion systems include grate burners, some fluidized-bed burners, and more recently gasification and pyrolysis. Heat generated in these combustion systems is used to raise steam and drive a steam generator.

Within the broad outline above, power-from-waste plants vary enormously. Much depends on the waste to be burned, its energy content, the amount of recyclable material or metal it contains, and its moisture content. Waste may be sorted before combustion or it may be burned as received. Emission control systems will vary too, with toxic metals and dioxins a particular target, but nitrogen oxide, sulfur dioxide, other acidic gases, and carbon monoxide emissions must all fall below local limits. Carbon dioxide emissions may need monitoring to comply with greenhouse gas emission regulations.

Once the waste has been burned, residues remain. Power-from-waste plants will generally reduce the volume of waste to around 10% of its original. A way must then be found to dispose of this residual ash. If it is sufficiently benign, it may be used as aggregate for road construction. Otherwise, it will probably be buried in a landfill. Other residues from emission control systems will have to be buried in controlled landfill sites too.

Northern Europe has been the traditional home of waste incineration plants for power generation and it continues to house the largest concentration of such plants. Altogether there are around 440 WTE plants in the EU producing 30 TWh of electricity and 55 TWh of heat in 2009.[4] Japan has also made extensive use of waste combustion, though not always for power generation, with around 100 plants in operation, while the United States has a similar number. In 2011 there were about 800 WTE plants in operation in 40 countries around the world. These plants were estimated to have treated 11% of MSW generated globally.

Europe has also developed the most widely used waste combustion technology based on waste incineration. Two companies, Martin GmbH based in Munich, Germany, and Swiss company Von Roll, accounted for close to 70% of the market for the dominant technology called mass-burn at the end of the 20th century.[5] The rest of the market is divided among a number of smaller companies, mostly based in either Europe, the United States, or Japan. The dominant European technology has also been widely licensed. It was the source

4. Figures are from the Confederation of European Waste to Energy Plants.

5. N. J. Themelis, An Overview of the Global Waste-to-Energy Industry, *Waste Management World*, July–August 2003.

of the technology used in most U.S. power-from-waste plants built in the late 1970s and early 1980s. More recently several developing countries of Asia have taken interest in power-from-waste and European technology has been modified for use in China, for example.

At the same time newer technologies based on gasification and pyrolysis are being developed by a variety of companies. These are based on technologies from other industries such as petrochemicals.

TRADITIONAL WASTE INCINERATION PLANTS

The traditional method of converting waste to energy is by burning it directly in a special combustion chamber and grate, a process that is often called mass burning (Figure 16.2). The dominant European technologies use this system. These involve specially developed moving grates, often inclined to control the transfer of the waste, and long combustion times to ensure that the waste is completely destroyed. Designs have evolved over 20–30 years and are generally conservative.

More recently, fluidized-bed combustion systems have sometimes been used in place of traditional grates. Such systems are good at burning heterogeneous fuel but require it to be reduced to small particles first.

The actual grate, be it conventional or fluidized bed, forms only a part of a waste treatment plant. A typical solid waste combustion facility is integrated into a waste collection infrastructure. Waste is delivered by the collecting trucks to a handling (and possibly a sorting) facility where it must be stored in a controlled environment to prevent pollution. Recyclable materials may be removed at this stage, though metallic material may be recovered after combustion. Grabs and conveyors will then be used to transfer the combustible waste from the store to the combustor.

Plant components, and particularly the grates, must be made of special corrosion-resistant materials. The grate must also include a sophisticated combustion control system to ensure steady and reliable combustion, while the quality and energy content of the refuse fuel varies. In some more modern systems oxygen is fed into the grate to help control combustion. The temperature at which the combustion takes place must usually be above 1000 °C to destroy chemicals such as dioxins, but must not exceed 1300 °C as this can affect the way ash is formed and its content.

Hot combustion gases from the grate flow vertically into a boiler where the heat is captured to generate steam. The combustion process in the grate and the temperature profiles within the boiler have to be maintained carefully to control the destruction of toxic chemicals. Most of the residual material after combustion is removed from the bottom of the combustion chamber as slag. However, there may be further solid particles in the flue gases, some of which can be recycled into the furnace.

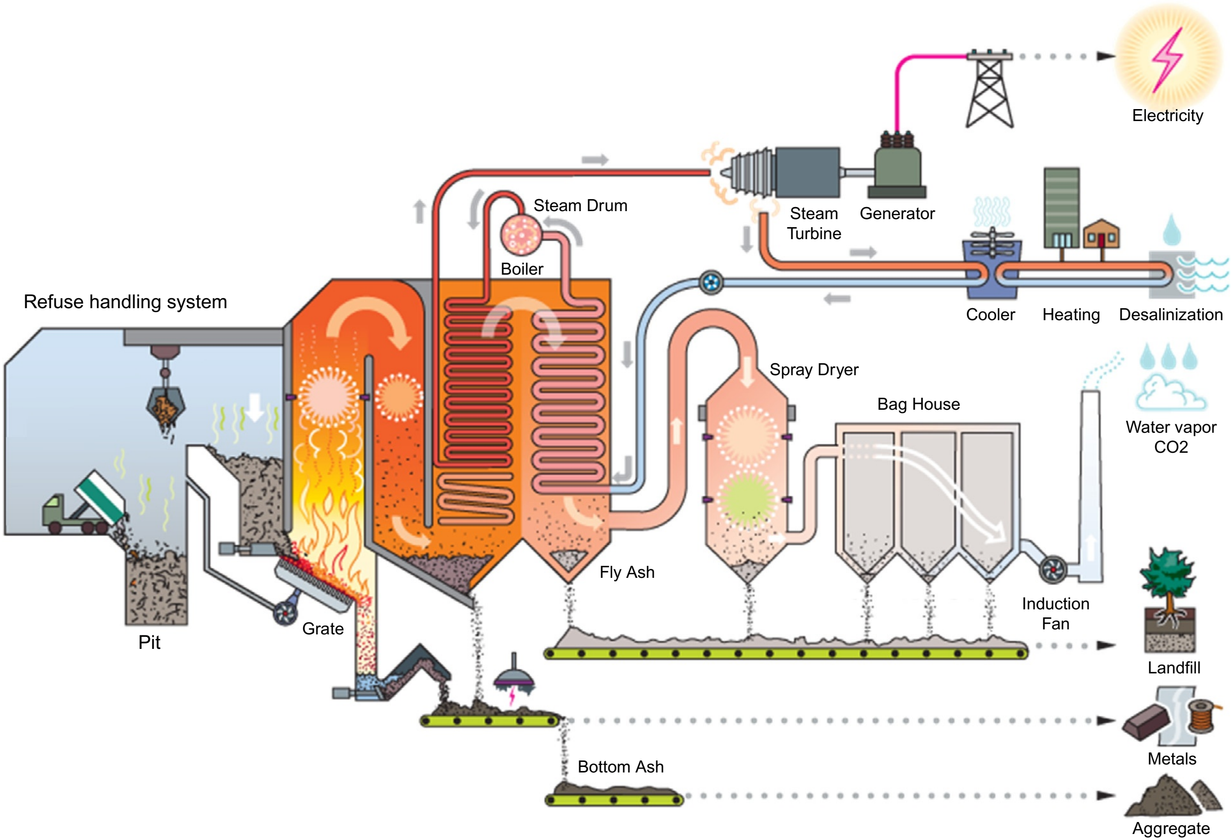

FIGURE 16.2 Schematic of an urban mass-burn power from waste plant.

Upon exiting the combustion and boiler system, the exhaust gases have to be treated extensively. While the combustion chamber may utilize techniques to minimize nitrogen oxide emissions—though further reduction may prove necessary—a system to capture sulfur will be required. This will probably be designed to capture other acidic gases such as hydrogen-chloride too. There may be a further capture system based on active carbon that will absorb a variety of metallic and organic residues in the flue gases. Then some sort of particle filter will be needed to remove solids. By this stage the exhaust gases should be sufficiently clean to release into the atmosphere, but continuous monitoring systems are required to make sure emission standards are maintained.

Dust from the flue-gas filters is normally toxic and must be disposed of in a landfill. Other flue-gas treatment residues will probably need to be buried too. The slag from the combustor may, however, be clean enough to exploit for road construction. Modern mass burn plants aim to generate slag that can be utilized in this way.

Mass burn plants may burn up to 2000 tonnes/day of MSW. Where a smaller capacity is required, a different type of combustion system, called a rotary kiln, can be employed. As its name suggests, this system uses a rotating combustion chamber that ensures that all the waste is burned. The chamber is inclined so that the material rolls from one end to the other as it burns. Such combustors are capable of burning waste with a high moisture content, perhaps up to 65%. Capacities of rotary kilns are up to 200 tonnes/day of refuse, suitable to meet the needs of small urban communities.

GASIFICATION AND PYROLYSIS

In recent years a number of companies have developed new WTE technologies based on both gasification and pyrolysis. These technologies are derived from the power and the petrochemicals industries. Pyrolysis is a partial combustion process carried out at moderate to high temperatures in the absence of oxygen and it can produce a mixture of gaseous, liquid, and solid residues. The traditional method of producing charcoal is a form of pyrolysis, as is the production of gas from coal which leaves a residue of coke. Gasification, meanwhile, involves heating solid material at high temperatures in a limited amount of air or oxygen to produce a char waste and a combustible gas. In both cases the gas will normally be burned to generate heat and steam.

Pyrolysis can produce a range of products from waste, depending on the temperature and the time the waste spends in the pyrolysis reactor. At lower temperatures with short residence times, more oils and tars are produced. Longer residence times lead to more solid residue (char). When the temperature is low, these solid and liquid products can contain complex and sometimes toxic organic molecules and must generally then be combusted at a controlled high temperature to generate power in a conventional boiler system.

Many WTE pyrolysis plants operate at relatively high temperatures so that they do not produce any liquid or tar residues, only combustible gas and a solid residue. Waste is normally sorted first, removing and recycling as much metal, glass, and plastics as possible. The remainder is then shredded and reduced to small particles of perhaps 2 mm in size before exposing it to a high temperature of around 800 ºC to convert it very quickly into a combustible gas and solid ash. The gas, which has a calorific value up to 22 MJ/kg, is cleaned and can then be used in a gas engine to generate power. Alternatively, the gas can be used in a conventional boiler.

One of the main advantages touted for high-temperature pyrolysis over incineration is that production of toxic organic compounds, such as dioxins and furans, is minimized. One company's pyrolysis WTE plant is designed to treat around 90 tonnes a day of waste with a calorific value of 8.4 MJ/kg and a moisture content of 40%, generating 2 MW of power in the process.[6]

Another, lower-temperature pyrolysis system, developed in the 1990s in Japan,[7] employs an initial pyrolysis process followed by combustion to generate heat. Waste delivered to the plant is first shredded and then fed into a rotating pyrolysis drum where it is heated to around 450 ºC. The heat, provided by hot air generated at a later stage in the process, pyrolyses the waste, converting it into a combustible gas and a solid residue.

The solid residue contains any metal that entered with the waste. This can be removed at this stage for recycling. Both iron and aluminum can be segregated in this way. The remaining solid slag is crushed. The gas and the crushed residue are then fed into a high-temperature combustion chamber operating at 1300 °C where it is completely burned. Combustion is controlled to limit nitrogen oxide formation. Incombustible material adheres to the walls of the combustion chamber where it flows, in liquid form, to the bottom. From here it is led out of the bottom of the furnace and immediately quenched, creating an inert granular material suitable for road building.

Hot flue gases from the combustion chamber are used to generate steam to drive a turbine. Dust is then removed from the exhaust gases and returned to the combustion system. Following this, a flue-gas treatment system removes any remaining acid gases. Only this material, around 1% of the original volume of the MSW, needs to be disposed of in a landfill. This system also claims to keep residual levels of dioxins extremely low.

Waste gasification (Figure 16.3) is similar to pyrolysis but it generally takes place at a higher temperature, typically 1000 °C, to produce a combustible, low calorific–value gas that can be burned either in a gas engine or in a conventional boiler system. As with pyrolysis, the products of gasification depend on the temperature used and lower temperatures can lead to more contaminants in the gas. The

6. These figures are for WPP Energy Corp's pyrolysis WTE plant.

7. The process, called R21, was developed by Mitsui Engineering and Shipbuilding. The first plant was completed in 2000.

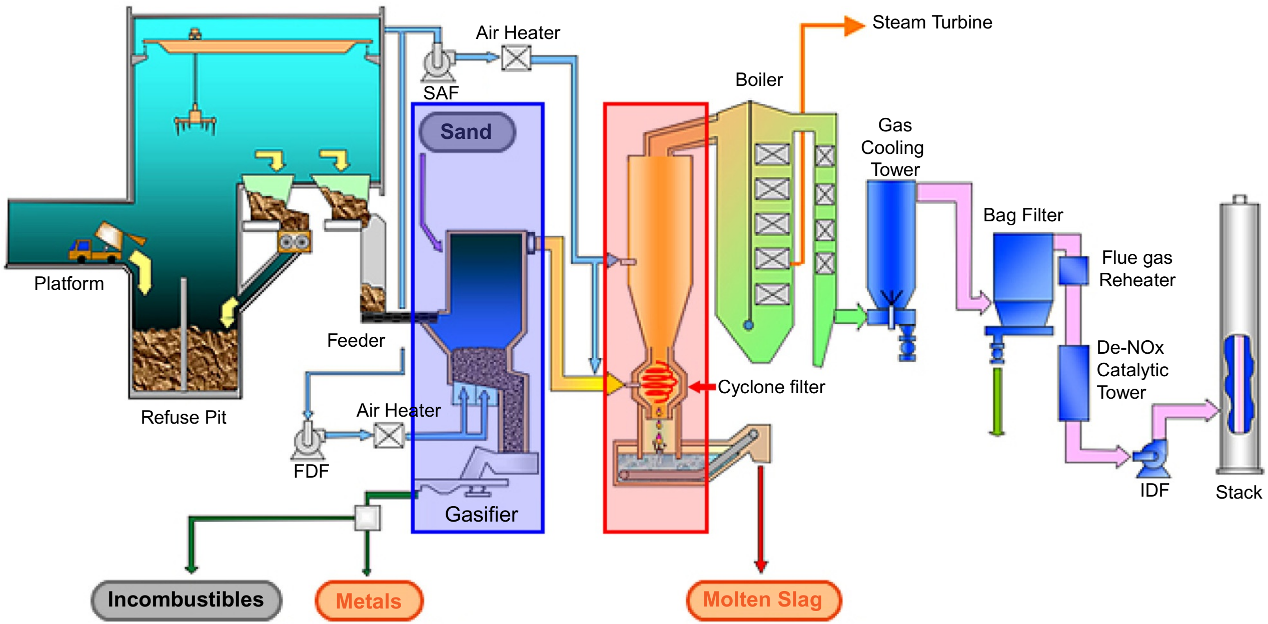

FIGURE 16.3 Schematic of an urban waste gasification plant.

synthetic gas produced during gasification is generally a mixture of carbon monoxide and hydrogen but the carbon monoxide can be converted into more hydrogen by using a shift reaction in which the gas is mixed with water vapor and passed over a catalyst at high temperatures. Provided it is clean enough, this gas can be used either for power generation or as a feedstock for industrial processes.

A new version of the gasification process that is being developed for MSW processing is plasma gasification (Figure 16.4). This involves burning the waste in a plasma arc at temperatures that can reach in excess of 10,000 °C, although the actual gasification will generally take place at slightly lower temperatures. Even so, the temperature is so high that all the components of the waste are broken down to atomic constituents and the product is expected to be a relatively pure syngas.

A plasma gasification reactor consists of a gasification chamber with a plasma torch generating a high-temperature arc close to the bottom. Solid waste is introduced above the torch, and as it falls through the high-temperature region, it is gasified in the presence of a controlled amount of air introduced from the bottom of the chamber. Any metal and slag is removed from the bottom of the reactor while syngas exits from the top. The high temperatures reduce the levels of chemicals such as dioxins to a very low level, often much lower than the levels produced in a conventional incineration plant.

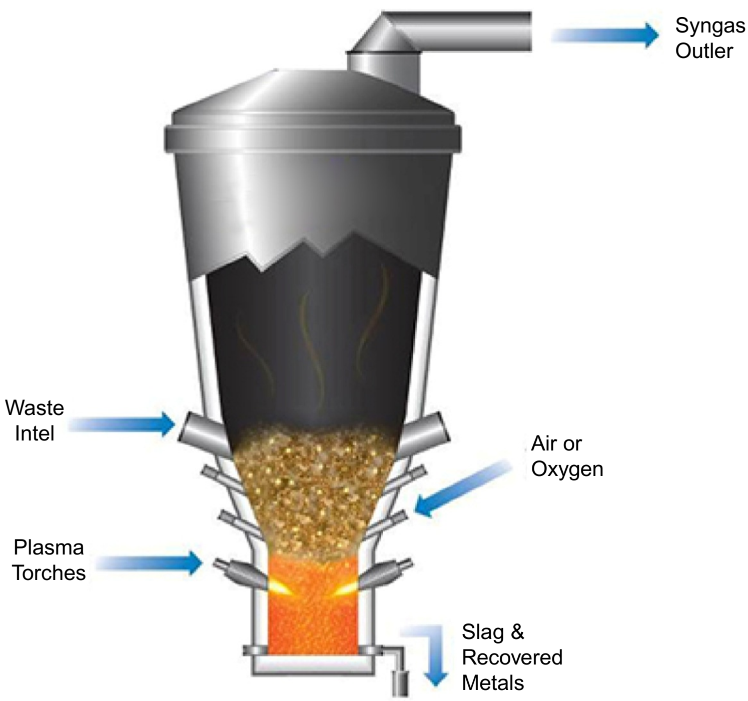

FIGURE 16.4 Cross-section of a plasma gasifier.

The syngas from a plasma gasifier can be burned in an engine or conventional combustion plant. It is also possible to feed the gas directly into a coal-fired power plant in the same way as was discussed for biomass gasification in Chapter 15.

REFUSE-DERIVED FUEL

RDF is the product of the treatment of MSW to create a fuel that can be burned easily in a combustion boiler. To produce RDF, waste must first be shredded and then carefully sorted to remove all noncombustible material such as glass, metal, and stone. Shredding and separating is carried out using a series of mechanical processes that are energy intensive. The World Bank has estimated that it requires 80–100 kWh to process one tonne of MSW and a further 110–130 kWh to dry the waste.[8]

After the waste has been shredded and separated, the combustible portion is formed into pellets that can be sold as fuel. The original intention of this process was to generate a fuel suitable for mixing with coal in coal-fired power plants. This, however, led to system problems and the modern strategy is to burn the fuel in specially designed power plants. An alternative is to mix the RDF with biomass waste and then burn the mixture in a power plant. Since RDF production must be preceded by careful sorting, this type of procedure is best suited to situations where extensive recycling is planned.

ENVIRONMENTAL ISSUES

Urban waste, its production, and its fate are major environmental issues. As already noted, modern urban living produces enormous quantities of waste in the form of paper, plastic, metals, and glass, as well as organic materials. How these wastes are processed is a matter of increasing global concern. Wastes such as paper, glass, and metal can be recycled, as—in theory—can plastics. From an environmental perspective it makes sense to reuse as much waste as possible, so environmentalists generally favor maximum recycling. Many governments now promote recycling too. However, the economics of recycling are not clear-cut and there are critics who consider it economically ineffective. Since such debates pitch sustainability against economics, the issue is not easily resolved.

While recycling offers the ideal solution, in practice there are frequently neither the facilities nor the infrastructure to recycle effectively. Even where recycling is employed there is still a residue of waste that cannot be reused. Thus, there remains a considerable volume of waste for which an alternative means of

8. Mining the Urban Waste Stream for Energy: Options, Technological Limitations, and Lessons from the Field, U.S. Agency for International Development, 1996 (Biomass Energy Systems and Technology Project, DHR-5737-A-00-9058-00).

disposal is required. The only options currently available are burial in a landfill site or combustion.

The combustion of waste would seem initially to be the ideal solution. Combustion reduces the quantity of waste to 10% or less of its original volume. At the same time it produces energy as a by-product and this energy can be used to generate electricity, heating, or both. Unfortunately, waste often contains traces of undesirable substances that may emerge into the atmosphere as a result of combustion. Other hazardous products may result from the combustion itself, with the waste providing the chemical precursors. So, while solving one environmental problem, waste combustion can generate others.

In the face of this, the combustion of waste has become subject to strict legislation. This sets limits on amounts of different hazardous materials that can be released as a result of the process. Chief among these are heavy metals, such as mercury, cadmium, and others, and potent organic compounds, such as dioxins and furans.

Modern WTE plants aim to meet the regulatory requirements imposed on them. In spite of this they are often extremely unpopular and even when the technology appears capable of limiting emissions to extremely low levels, it can be difficult gain permission to build plants. New waste conversion technologies such as gasification and pyrolysis may be able to help overcome popular objections if they can be shown to generate negligible toxic emissions.

WASTE PLANT EMISSIONS

A plant burning waste produces four major types of product. First, there is the solid residue from the grate itself, normally termed slag or ash. Second, there is the chemical product resulting from flue-gas treatment systems. Third, there is a quantity of dust in the flue gases emerging from the plant boiler; this is normally captured with filters or an electrostatic precipitator. Finally, there is the flue gas itself.

Ash

The nature of the ash or slag emerging from the grate of a power-from-waste plant will depend on both the type of waste being burned and the combustion conditions. While its primary constituents will be solid, incombustible mineral material from the wastes, this residue will be contaminated with traces of a variety of metals. These traces may be in a toxic or harmless form. If the waste has not been carefully sorted beforehand, the slag may also contain larger slugs of valuable metal that can be recycled.

By careful control of the temperature in the furnace, it is possible to incorporate the trace metals into the mineral content of the ash and render them effectively harmless. This is a process called sintering. The effectiveness of the sintering process in rendering toxic metals harmless will be determined by measuring the amounts capable of being leached out by water. The ash may also contain some toxic organic compounds such as dioxins. Furnace conditions can minimize these too since a sufficiently high temperature will normally destroy such compounds. The effectiveness of this will again be determined by a leaching test.

If the ash or slag is too toxic for any other use it will have to be buried in a landfill. Modern facilities aim to render it sufficiently stable and benign that it can be used for road building or similar purposes. When they succeed, only a residual 1% of the original waste needs to be buried.

Fly Ash and Flue-gas Treatment Residues

Fine, solid particles called fly ash escape with the flue gas from a furnace. This fly ash will often contain high levels of toxic metals and must be captured. Capture is achieved either by using a fabric filter called a bag filter, or by employing a device called an electrostatic precipitator. Both should be capable of removing close to 100% of the dust from the flue gas. Once captured this dust must be safely buried in a landfill.

Other treatment systems are designed to remove gaseous components from the flue gas. This includes acidic compounds and harmful organic compounds. The processes used to remove these components, often similar to the treatment plants used for coal-fired power plants and described in detail in Chapter 3, result in by-products that also require disposal. Depending on the treatment process, the residue may be a solid or wet slurry. In the latter case, the slurry will normally be dried using the hot exhaust gases before disposal.

Flue Gas

Once treated, the flue gas from a waste combustion plant should be sufficiently clean to release into the atmosphere. However, the gas will usually need to be monitored to ensure that emission limits are being met.

Dioxins

One of the most potent environmental concerns during the last 20–30 years has related to the release of dioxins into the atmosphere. Dioxins are undesirable by-products of the manufacture of a variety of chemicals such as pesticides and disinfectants, but one particular compound called 2,3,7,8-tetrachlorodibenzo-p-dioxin has come to be identified as dioxin. This material was thought to be extremely toxic to humans, though more recent studies suggest earlier results were exaggerated. However, several dioxins are considered carcinogenic.

Dioxins can be found in urban waste but the principal danger is that the compounds are formed during waste combustion if the process is not carefully controlled. This is usually a matter of temperature. If the combustion temperature is too low some plastics and other materials can break down and then the components react to create dioxins. Once the combustion temperature is high enough, usually above 1000 ºC, then these compounds will normally not be produced. Some early waste incineration plants did not control the emissions sufficiently carefully and this led to instances of widespread contamination. Such instances have colored the perception of WTE plants ever since.

Dioxin emission levels are now closely regulated and emissions have fallen. In the United States, according to the Environmental Protection Agency, the total emissions of dioxins from large WTE facilities fell from over 8 kg (toxic equivalent) in 1987 to less than 14 g (toxic equivalent) by 2005. The European emission limit for dioxins is 0.1 ng/Nm^3. Power-from-waste plants built in the first decade of the 21st century and beyond should be capable of reducing the emission level to one-tenth or even one-hundredth of this.

Heavy Metals

Heavy metals, particularly lead, cadmium and mercury have proved another source of concern. Less mercury is used today than in the past. This, combined with better filtration systems has reduced mercury emissions from power-from-waste plants in the USA from over 50 tonnes/year in 1990 to around 2 tonnes/year in 2005. Coal-fired power plants release over 40 tonnes/year. Lead and cadmium emissions have fallen by 96% over the same period. In general the emissions of toxic metals from waste incineration plants should fall well below legal emission limits.

COST OF ENERGY FROM WASTE PLANTS

The capital cost of equipment to generate electricity from waste is generally much higher than for conventional power generation equipment to burn fossil fuel. Plant design is specialized and must include refinements for emission control that are not necessary in a fossil fuel plant. Grate design is unique too.

Against this must be offset the revenue of the plant, not only from the electricity generated but also from the fuel itself, the waste. Industry and municipalities expect to pay to dispose of their waste. Consequently, the economics of a project should be designed so that the revenue from the waste disposal contracts are adequate to enable the power from the plant to be sold competitively. It should be remembered, however, that a waste to energy plant has as its primary purpose the treatment and elimination of MSW. Electricity is a useful by-product of this process but generation is not the main function of the plant.

A study carried out for the Mayor of London and published in 2008 looked at the cost of the principle waste combustion technologies. The main findings are shown in Table 16.2. The study concluded that a conventional incineration facility would cost around £45 m for a plant with the capacity to treat 100,000 tonne/y of MSW while for a 200,000 tonne/y plant the cost would

TABLE 16.2 Cost of WTE Plants in the United Kigdom

Plant waste treatment capacity	Conventional incineration	Advanced thermal treatment
100-115 ktonne/y	£45 m	£50 m
150 ktonne/y	£60 m	£68 m
170-200 ktonne/y	£76 m	£85 m

Source: Costs of Incineration and Non-incineration Energy-from-Waste Technologies, The Mayor of London, 2008.

be £76 m. With the maximum power output from the smaller plant put at 6 MW, this equates to £7500/kW while the larger plant has a maximum output of 12 MW, equating to 6300/kW.

Advanced thermal treatment plants such as gasifiers and pyrolysis plants have slightly higher costs, as shown in the table. Their potential power outputs are also slightly lower. As a consequence the unit cost of a 100,000 tonne/y advanced plant is £9100/kW while for the 200,000 tonne/y plant the unit cost is £7700/kW. Operating costs for the plants are broadly similar at between £40/tonne and £70/tonne depending upon plant size.

U.K. costs are similar to estimates for plants in the United States where the cost of a typical municipal waste combustion plant was put at $5,000–10,000/kW during the middle of the first decade of the 21st century. Again smaller plants are relatively more expensive than larger plants.

Chapter 17

Nuclear Power

Nuclear power is the most controversial of all the forms of power generation. To evaluate its significance involves weighing political, strategic, environmental, economic, and emotional factors that attract partisan views far more strident that any other method of electricity generation.

There are two ways of generating energy from nuclear reactions: nuclear fission and nuclear fusion. Both have their roots in the atomic weapons programs of World War II, but only nuclear fission has developed to be capable of commercial deployment in power plants and it is this form that has the highest profile. Commercial power generation based on fusion remains at least 20 years away.

Fission was at the heart of the first atom bomb, and work on the development of nuclear fission as a source of electricity gathered momentum during the late 1940s in both the United States and Russia. It was in the United States that an experimental breeder reactor at the Argonne National Laboratory, which started operating in 1951, first produced a small amount of electricity. Meanwhile, in Russia a water-cooled, graphite-moderated reactor called AM-1 was the first nuclear generating plant. It had a capacity of 5 MW when it began operation in 1954. In 1956 two 65 MW dual-purpose reactors started at Calder Hall in Cumbria, U.K., and in 1957 the U.S. Atomic Energy Commission built the 60 MW Shippingport pressurized water reactor, the first demonstration for a commercial nuclear reactor. Russian development of commercial power plants lagged behind that of the United States and it was not until 1964 that the first 210 MW plant at Novovoronezh entered service.

From these beginnings, nuclear power grew rapidly so that by the beginning of the 1970s it had blossomed into the great hope for unlimited global power. In 1974, the U.S. power industry alone had ordered 200 nuclear reactors, and in the same year the U.S. Energy Research and Development Administration estimated that U.S. nuclear-generating capacity could reach 1200 GW by 2000. (In fact, total U.S. generating capacity in 2000 from all sources was less than 1000 GW.[1]) The United Kingdom, France, Germany, Soviet Union, and Japan had all begun to build up substantial nuclear-generating capacities too.

1. U.S. Department of Energy.

But even as orders were being placed, the nuclear industry was reaching a watershed. A combination of economic, regulatory, and environmental factors were about to conspire to bring the development of nuclear power to a halt in the United States and across most of the developed world. There were already environmental and safety concerns during the 1970s, but two accidents—one at Three Mile Island in the United States in 1979 and a second at Chernobyl in the Ukraine in 1986—turned public opinion strongly against nuclear power. In response new safety regulations were introduced in the United States, lengthening construction times and increasing costs. As a consequence of this, and of government decisions elsewhere, nuclear construction across the world almost came to a standstill. By the late 1980s, a 100 nuclear projects in the United States had been canceled. To make matters worse the question of how to dispose of nuclear waste became a political issue.

The United States still retains a large fleet of nuclear power stations that it continues to operate, and work has recently started on new reactors after a hiatus of many years. However, some countries in Europe and Scandinavia decided to rule out the nuclear option completely. In 1978 Austria voted to ban nuclear power. Sweden voted in 1980 to phase out nuclear power by 2010, although this policy was repealed in 2010. Italy closed its last reactors in 1990. Germany reached an agreement with its nuclear power producers in 2000 to phase out its nuclear stations. Other western countries such as France, Belgium, and Finland remain positive about nuclear generation. The U.K. government, too, retains a nuclear option. And in 2003 the Finnish utility TVO ordered a new nuclear unit, the first to be built in the EU for over a decade.

There remained a large fleet of nuclear power plants in eastern Europe too. These plants are all based on Russian-designed reactors. The safety of the Russian designs was a matter of concern after the Chernobyl accident in 1986, and from the beginning of the 1990s, when Cold War barriers fell, efforts were made to improve the safety of eastern European reactors or to force their closure.

The evolution of nuclear generation in Asia followed a different course. Japan continued to develop its installed nuclear base, as did South Korea, though the Japanese nuclear industry began to face considerable public criticism at the end of the 20th century. Taiwan ordered two new nuclear reactors in 1996 but these had still not been completed at the end of 2012, and they are a source of controversy in the county. India has an indigenous nuclear industry. And in the mid-1990s, China started to develop what has now become a strong nuclear base.

By the end of the first decade of the 21st century there was hope within the nuclear industry that it was about to see a renaissance in the west too. New, safer plants had been developed and several western nations were considering new projects. However in March 2011, the nuclear industry was rocked by a further disaster when an earthquake and tsunami off the coast of Japan had a devastating impact on the four reactors at the Fukushima Daiichi nuclear power plant on the Japanese coast. The tsunami damaged the cooling systems of three reactors, and over a three-day period their cores largely melted, releasing radioactive material into the atmosphere.

This new nuclear disaster has led to more questioning of nuclear power and the hoped-for renaissance may have been stopped before it started. Japan's politicians are now grappling with the legacy and with its implication for the country's remaining nuclear plants. As a result of this accident, the German government has advanced the shutdown of all its nuclear power plants and closed all those that began operating before 1980. The remainder will now be finally retired by 2022 instead of 2036. The Chinese government halted nuclear power plant construction and placed a moratorium on the approval of new nuclear projects, although that was lifted in late 2012. The impact elsewhere is yet to be clear but French and U.K. national programs could face a more difficult future too.

While nuclear fission reactors continue to excite controversy and debate, some technologists see nuclear fusion as the safer long-term alternative. Nuclear fusion, the basis for the hydrogen bomb, has yet to develop to the stage when it can be used to generate electrical power. However, fission scientists and technologists continue to make advances. Fusion is inherently safer than fission but it remains a nuclear technology, and if it ever does establish itself as commercially viable, then it will have to prove itself against the background of the nuclear fission industry.

GLOBAL NUCLEAR CAPACITY

At the end of 2009, according to figures compiled by the World Energy Council[2] there were 437 operating nuclear reactors worldwide. These had a total generating capacity of 370 GW. A further 55 units were under construction; these had an aggregate capacity of 51 GW.

The global figures are broken down in Table 17.1 to show the distribution of current nuclear-generating capacity by region. Europe, with 195 units and 168 GW, has the greatest capacity. North America has 124 operating units with an aggregate generating capacity of 115 GW, while Asia has 112 units. Of the continents, only Australia and Antarctica have none.

Nationally, France produces around 75% of its electricity from nuclear power stations. Lithuania generates 76% from nuclear sources and Belgium 52%. In Asia, South Korea produces 35% of its power from nuclear units, while Japan before the Fukushima disaster, relied on nuclear power for 29% of its electricity. In all, 15 countries rely on nuclear plants for 25% or more of their electricity.

Net global nuclear generation in 2009 was 2558 TWh. Globally, nuclear power plants provide around 14% of total electricity generation, almost as much as hydropower. However, the overall global nuclear capacity has become

2. 2010 Survey of Energy Resources, World Energy Council.

TABLE 17.1 Global Nuclear-generating Capacity

	Number of Units	Total Capacity (MW)
Africa	2	1800
North America	124	114,560
South America	4	2701
Asia	112	82,642
Europe	195	168,484
Middle East	1	915

Source: 2010 Survey of Energy Resources, World Energy Council.

relatively static with new plants built in Asian countries such as China compensating for old plants removed from service in other parts of the world. Economically, nuclear power plants are perceived to be expensive to build. However, plants where the capital cost has been written off have proved extremely competitive generators of electricity, particularly in the United States.

How global capacity will develop in the future has become uncertain after the Fukushima Daiichi accident. The major source of nuclear growth over the past decade has been China and this continues to be the case, but there are indications that ambitions here may be more limited than previously. Japan's nuclear industry has suffered a damaging blow that may eventually force the country to retreat from nuclear fission. In Europe both France and the U.K. governments are supportive of nuclear power, but public support may be waning and the economics of construction in both countries are being questioned. Meanwhile the United States, which still has the largest nuclear fleet, in the world has taken the first steps in the second decade of the 21st century to build new nuclear capacity.

FUNDAMENTALS OF NUCLEAR POWER

All the nuclear power stations operating today generate electricity by utilizing energy released when the nuclei of a large atom such as uranium split into smaller components, a process called nuclear fission. The amount of energy released by this fission process is enormous. One kilogram of naturally occurring uranium could, in theory, release around 140 GWh of energy. (140 GWh represents the output of a 1000 MW coal-fired plant operating at full power for nearly six days.)

There is another source of nuclear energy, nuclear fusion, which involves the reverse of a fission reaction. In this case, small atoms are encouraged to fuse at extraordinarily high temperatures to form larger atoms. Like nuclear fission, fusion releases massive amounts of energy. However, it will only take place under extreme conditions. The fusion of hydrogen atoms is the main source of energy within the sun.

The reason why both fission and fusion can release energy lies in the relative stability of different elements. It turns out that atomic species in the middle of the periodic table of elements, elements such as barium and krypton (these are typical products of uranium fission), are generally more stable than either lighter elements such as hydrogen or heavier elements such as uranium. The nuclei of these more stable atoms are bound together more strongly and their nuclear components, the protons and the neutrons, are in fact slightly lighter. It is this difference in mass, equivalent to the stronger binding, that is released during the fission or fusion reaction.

Nuclear Fission

Many large, and even some small, atoms undergo nuclear fission reactions naturally. One of the isotopes of carbon (isotopes are atoms of a single element with different numbers of neutrons) called carbon-14 behaves in this way. Carbon-14 exists at a constant concentration in natural sources of carbon. Thus, living entities that constantly exchange their carbon with the biosphere maintain this constant concentration. However, when they die, the carbon-14 is no longer renewed and it gradually decays. Measuring the residual concentration gives a good estimate of the time since the organism died. It is this property that allows archeologists to use carbon-14 to date ancient artifacts and remains.

Other atoms can be induced to undergo fission by bombarding them with subatomic particles. One of the isotopes of uranium, the element most widely used in nuclear reactors, behaves in this manner. Naturally occurring uranium is composed primarily of two slightly different isotopes called uranium-235 and uranium-238 (the numbers refer to the sum of protons and neutrons each atom contains). Most uranium is uranium-238, but 0.7% is uranium-235.

When an atom of uranium-235 is struck by a neutron it may be induced to undergo a nuclear fission reaction. The most frequent products of this reaction are an atom of krypton, an atom of barium, three more neutrons, and a significant quantity of energy:

$$^{237}_{92}\mathrm{U} + n = ^{140}_{56}\mathrm{Ba} + ^{96}_{36}\mathrm{Kr} + 3n + @200\,\mathrm{meV}$$

In theory, each of the three neutrons produced during this reaction could cause three more atoms of uranium-235 to split. However, this also depends on the quantity of uranium present. If a piece of uranium is too small, then most of the neutrons will escape into the surroundings without ever meeting a uranium nucleus. It is only when the size reaches and exceeds a quantity known as the critical mass that the number of reactions created by each single fission reaction exceeds one. This leads to a rapidly accelerating reaction, called a chain reaction, which will release an enormous amount of energy. A chain reaction of this type forms the basis for the atomic bomb.

In fact, a lump of natural uranium will not explode because the uranium-235 atoms only react when struck by slow-moving neutrons. The neutrons created during the fission process move too fast to cause further fission reactions to take place. They need to be slowed down first. This is crucial to the development of nuclear power.

Controlled Nuclear Reaction

If uranium fission is to be harnessed in a power station, the nuclear chain reaction must first be tamed. The chain reaction is explosive and dangerous. However, it can be managed by carrying away the energy released by the fission reactions, controlling the number of neutron within the reactor core, and then slowing the remaining neutrons so that they can initiate more fission reactions.

An accelerating chain reaction will take place when each fission reaction causes more than one further identical reaction. If the fission of a single uranium-235 atom causes only one identical reaction to take place, the reaction will carry on indefinitely—or at least until the supply of uranium-235 has been used up—without accelerating. But if each fission reaction leads to an average of less than one further reaction, the process will eventually die away naturally.

The operation of a nuclear reactor is based on the idea that a nuclear chain reaction can be controlled so that the process will continue indefinitely but will never run away and become a chain reaction. A reactor in which each nuclear reaction produces one further nuclear reaction is described as critical. Once the product of each nuclear reaction is more than one additional reaction, the reactor is described as supercritical. Operation must be controlled so that the reactor is just—but barely—supercritical.

Naturally occurring uranium can be used as fuel for a nuclear fission reactor. However, most nuclear reactors contain uranium that has been enriched so that it contains more uranium-235 than it would in nature. Enrichment to about 3% is common. Using enriched uranium makes it easier to start a sustained nuclear fission reaction.

In addition to the uranium, the reactor also contains rods made of boron. Boron is capable of absorbing the neutrons generated during the nuclear reaction of uranium-235. If a sufficiently large amount of boron is included within the reactor core it will absorb and remove the neutrons generated during the fission reaction, stopping the chain reaction from proceeding by keeping the reactor subcritical. The boron rods are moveable, and by moving the rods in and out of the reactor core, the number (or flux) of neutrons and therefore the nuclear process can be controlled.

One further crucial component is needed to make the reactor work—something to slow down the fast neutrons. The neutrons from each uranium-235 fission move too fast to stimulate a further reaction, but they can be slowed by adding a material called a moderator. Water makes a good moderator and is used in most operating reactors. Graphite also functions well as a moderator and has been used in some reactor designs.

When a uranium fission reaction takes place the energy it releases emerges as kinetic energy. In other words, the products of the fission process carry the energy away as energy of motion; they move extremely fast. Much of this energy is carried away by the fast neutrons. These neutrons will dissipate their energy in collision with atoms and molecules within the reactor core. In many reactors this energy is absorbed by the moderator: water. So while the neutrons are slowed, the water within the core becomes hotter. By cycling the water from the reactor core through a heat exchanger, this heat can be extracted and used to generate electricity. Extracting the heat also helps maintain the reactor in a stable condition by preventing overheating.

The operation of a nuclear fission reactor is, therefore, a careful balancing act. As a consequence, a reactor always has the potential to generate a runaway chain reaction. Modern reactor designs try to ensure that there is no possibility of this happening in the event of a component or operational failure.

Nuclear Fusion

The alternative energy-yielding nuclear reaction to fission is fusion. Fusion is the process that generates energy in the sun and stars. In the sun, hydrogen atoms combine to produce deuterium (heavy hydrogen) atoms and then deuterium and hydrogen atoms fuse to produce helium with the release of energy. The reaction takes place at 10–15 million °C and at enormous pressure.

The conditions in the sun cannot be easily recreated on Earth, although fusion of the type taking place within the sun has been achieved. However, for the purposes of a fusion reactor capable of electricity generation, another reaction offers more potential because it takes place under more benign conditions than those in the sun. This is the reaction between two isotopes of hydrogen, deuterium, and tritium. Deuterium ${}^{2}_{1}H$ is found naturally in small quantities in water while tritium ${}^{3}_{1}H$ is made from lithium. These two will react to produce helium and energy:

$$ {}^{2}_{1}H + {}^{3}_{1}H = {}^{4}_{2}He + n + @18\,\text{meV} $$

The reaction between deuterium and tritium will only take place at 100 million °C (but at much lower pressure than in the sun). At this temperature all the atoms separate into a sea of nuclei and electrons, a state called a plasma. Since the constituents of a plasma are all charged, either positively or negatively, they can be controlled and contained using a magnetic field. This is crucial since there is no material that can withstand temperatures this severe. The most promising magnetic field for containing a plasma is torroidal and this has formed the basis for most fusion research. There is an alternative method of containing a fusion plasma called inertial confinement. This relies on generating extreme conditions within a small charge of tritium and deuterium, in essence creating a tiny sun in which the fusion takes place too fast for the particles to escape. Both systems of containment are being developed for power generation.

NUCLEAR FISSION REACTOR DESIGNS

Nuclear reactor is the name given to the device or structure in which a controlled nuclear reaction takes place. There are a number of different designs but these have many features in common.

The core of the reactor is its heart, the place where the nuclear fuel is placed and where the nuclear reaction takes place. The fuel is most frequently formed into pellets roughly 2 cm in diameter and 1–2 cm long. These pellets are loaded into a fuel rod, a hollow tube of a special corrosion-resistant metal; this is frequently a zirconium alloy. Each fuel rod is 3–4 m long and will contain 150–200 pellets. A single reactor core may contain up to 75,000 such rods. Fuel rods must be replaced once the fissile uranium-235 they contain has been used up. This is a lengthy process that can take as much as three weeks to complete during which the reactor normally has to be shut down. However, some designs allow refueling while in operation.

In between the fuel rods there are control rods, made of boron, that are used to control the nuclear reaction. These rods can be moved in and out of the core and they will be of different types. Some will be designed to completely stop the reaction in the core, others to adjust the speed of the reaction. The core will also contain a moderator to slow the neutrons released by the fission of uranium atoms. In some cases the moderator is also the coolant used to carry heat away from the core.

The outside of the core may be surrounded by a material that acts as a reflector to return some of the neutrons escaping from the core. This helps maintain a uniform power density within the core and allows smaller cores to be built. There may also be a similar reflecting material in the center of the core.

The coolant collects heat within the core and transfers it to an external heat exchanger where it can be exploited to raise steam to drive a steam turbine. The coolant may be water (light water), deuterium (heavy water), a gas such as helium or carbon dioxide, or a metal such as sodium. The core and its ancillary equipment are normally called the “nuclear island” of a nuclear power plant, while the boiler, steam turbine, and generator are called the “conventional island.” The coolant system will link the nuclear and conventional islands.

A nuclear power plant will contain a host of systems to ensure that the plant remains safe and can never release radioactive material into the environment. The most important of these is the containment. This is a heavy concrete and steel jacket that completely surrounds the nuclear reactor. In the event of a core failure the containment should be able to completely isolate the core from the surroundings and remained sealed, whatever happens within the core.

Boiling Water Reactor

The boiling water reactor (BWR) uses ordinary water (light water) as both its coolant and its moderator (Figure 17.1). In the boiling water reactor the water in the reactor core is permitted to boil under a pressure of 75 atmospheres, raising the boiling point to 285 °C, and the steam generated is used directly to drive a steam turbine. This steam is then condensed and recycled back to the reactor core. Since the steam is exposed to the core there is some radioactive contamination of the turbines, but this is short-lived and turbines can normally be accessed soon after shutdown.

This arrangement represents probably the simplest possible for a nuclear reactor because no additional steam generators are required. However, the internal systems within a BWR are complex. Steam pressure and temperature are low compared to a modern coal-fired power plant and the steam turbine is generally very large. BWRs have capacities up to 1400 MW and an efficiency of around 33%.

The BWR uses enriched uranium as its fuel. This fuel is placed into the reactor in the form of uranium-oxide pellets in zirconium-alloy tubes. There may be as much as 140 tonnes of fuel in 75,000 fuel rods. Refueling a BWR involves removing the top of the reactor. The core itself is kept under water, with the water shielding operators from radioactivity. Boron control rods enter the core from beneath the reactor.

In common with all reactors, the fuel rods removed from a BWR reactor core are extremely radioactive and continue to produce energy for some years. They are normally kept in a carefully controlled storage pool at the plant before, in principle at least, being shipped for either reprocessing or final storage.

The BWR was developed at the Argonne National Laboratory in the United States and a commercial version was subsequently produced by General

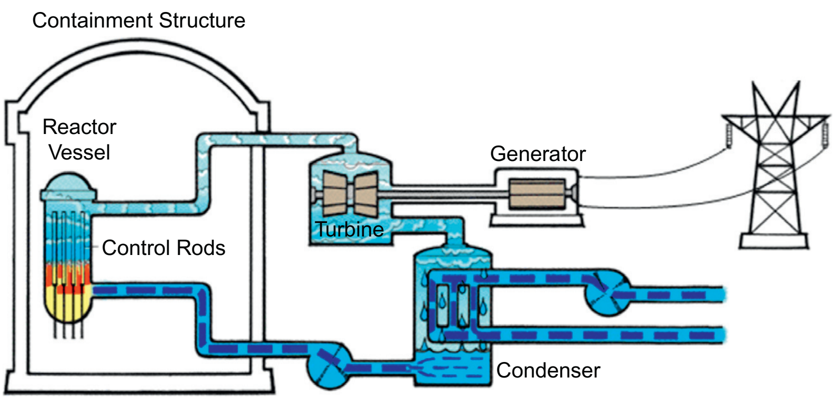

FIGURE 17.1 Schematic of a boiling water reactor.

Electric Co. (GE). There are close to 100 BWRs in operation including four advanced BWRs in Japan and two in Taiwan.

Pressurized Water Reactor

The pressurized water reactor (PWR) also uses ordinary or light water as both coolant and moderator (Figure 17.2). However, in the PWR system the cooling water is kept under pressure so that it cannot boil. The PWR differs in another respect from the boiling water reactor; the primary coolant does not drive the steam turbine. Instead, heat from the primary water cooling system is captured in a heat exchanger and transferred to water in a secondary system. It is the water in this second system that is allowed to boil and generate steam to drive the turbine.

The core of a PWR is filled with water, pressurized to 150 atmospheres, allowing the water to reach 325 °C without boiling. The use of a second water cycle introduces energy losses that make the PWR less efficient at converting the energy from the nuclear reaction into electricity. However, the arrangement has other advantages regarding fuel utilization and power density, making it competitive with the BWR. It also allows the reactor to be more compact.

The PWR uses enriched uranium fuel with a slightly higher enrichment level than in a BWR. This is responsible for a higher power density within the reactor core. As with the BWR, the fuel is introduced into the core in the form of uranium-oxide pellets. A typical PWR will contain 100 tonnes of uranium. Refueling is carried out by removing the top of the core. However, in a PWR the control rods are inserted from above, allowing gravity to act as a fail-safe in the event of an accident.

A typical PWR has a generating capacity of 1000 MW. The efficiency is around 33%. The PWR is the most popular reactor in use globally, with over 250 in operation. The most important commercial PWR was developed by

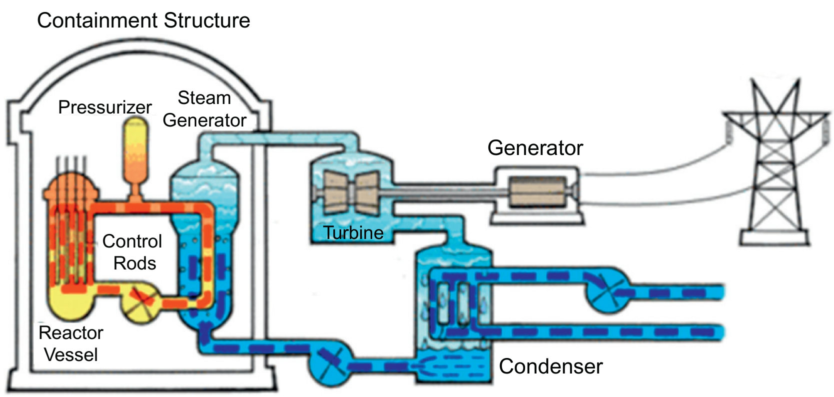

FIGURE 17.2 Schematic of a pressurized water reactor.

Westinghouse for ship propulsion and later converted to power generation. The Russians developed their own version of the PWR called the VVER and units of this type continue to operate in Russia and former Soviet countries. France also developed a PWR that was based on the Westinghouse design but the designs later diverged so that the French one is now an independent design.

Pressurized Heavy Water Reactor (CANDU Reactor)

The Canadian deuterium uranium (CANDU) reactor was developed in Canada with the strategic aim of enabling nuclear power to be exploited without the need for imported enriched uranium (Figure 17.3). Uranium enrichment is an expensive and highly technical process. If it can be avoided, countries such as Canada with natural uranium reserves can more easily exploit their indigenous reserves to generate energy. This has made the CANDU reactor, which uses unenriched uranium, attractive outside Canada too.

The CANDU reactor uses, as its moderator and coolant, a type of water called heavy water. Heavy water is a form of water in which the two normal hydrogen atoms have been replaced with two of the isotopic form—deuterium. Each deuterium atom weighs twice as much as a normal hydrogen atom, hence the name heavy water. Heavy water occurs in small quantities in natural water.[3]

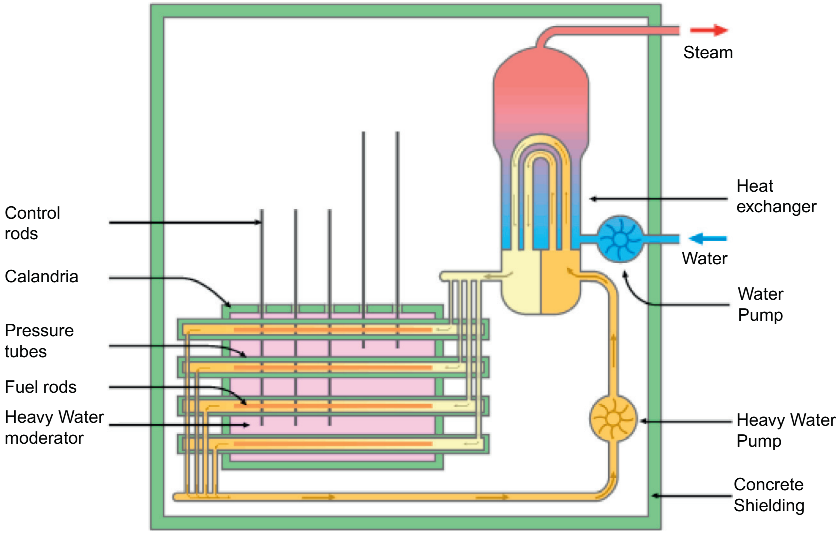

FIGURE 17.3 Schematic of a CANU reactor.

---

3. A deuterium atom is a hydrogen atom with an extra neutron, giving it twice the mass of normal hydrogen. About 1 in every 6760 naturally occurring hydrogen atoms is a deuterium atom. The two can be separated using electrolysis that selectively splits normal water but leaves heavy water.

Heavy water is much more expensive than light water but it has the advantage that it absorbs fewer neutrons than normal water. As a consequence, it is possible to sustain a nuclear reaction without the need to enrich the uranium fuel. The CANDU reactor has the additional advantage that it can be refueled without the need to shut it down; in fact, this is necessary with natural uranium fuel to keep the plant going. Avoiding lengthy refueling shutdowns provides better operational performance.

The CANDU fuel is loaded in the form of uranium oxide pellets housed in zirconium-alloy rods that are inserted horizontally into pressure tubes penetrating the core instead of vertically as in other PWRs and BWRs. Fuel replacement involves pushing a new rod into a pressure tube that passes through the vessel containing the heavy water (called a calandria) and forcing the old tube out of the other end. Since the pressure tube is isolated from the heavy water, refueling can be carried out without the need to shut down the reactor.

The heavy water coolant in the CANDU reactor is maintained under a pressure of around 100 atm and the water reaches around 290 °C without boiling. Heat is transferred through a heat exchanger to a light water system with a steam generator and the secondary system drives a steam turbine in much the same way as a PWR. Efficiency is similar too.

The CANDU reactor was developed by Atomic Energy of Canada, and that country has the largest CANDU fleet, but reactors have also been supplied to countries such as Argentina, South Korea, India, and Pakistan. There are around 40 in operation.

Gas-cooled Reactors

When searching for their own design of reactor that did not require enriched uranium, U.K. scientists developed the Magnox reactor. This reactor uses a graphite moderator and carbon dioxide as the heat-transfer medium. The latter carries away heat generated in the moderator to a heat exchanger where it is used to heat water and raise steam to drive a turbine. Channels in the graphite moderator contain tubes made of magnesium alloy (from which the reactor takes its name) into which uranium fuel is loaded.

The United Kingdom built nine Magnox reactors, all of which were different, and units were also built in Japan and Italy. A second generation of gas-cooled reactors, called advanced gas-cooled reactors (AGRs), were developed in the Untied Kingdom during the 1960s and seven of these were built (Figure 17.4). These retained the graphite moderator and carbon dioxide coolant but opted for 2% enriched uranium fuel housed in zirconium-alloy rods. The gas coolant in the AGR reaches 650 °C and the gas is then circulated through steam generator tubes outside the core.

The advantage of the AGR is that higher temperatures in the core can potentially provide for a higher efficiency of power generation. However, the U.K. fleet of gas-cooled reactors has not proved as successful, operationally, as

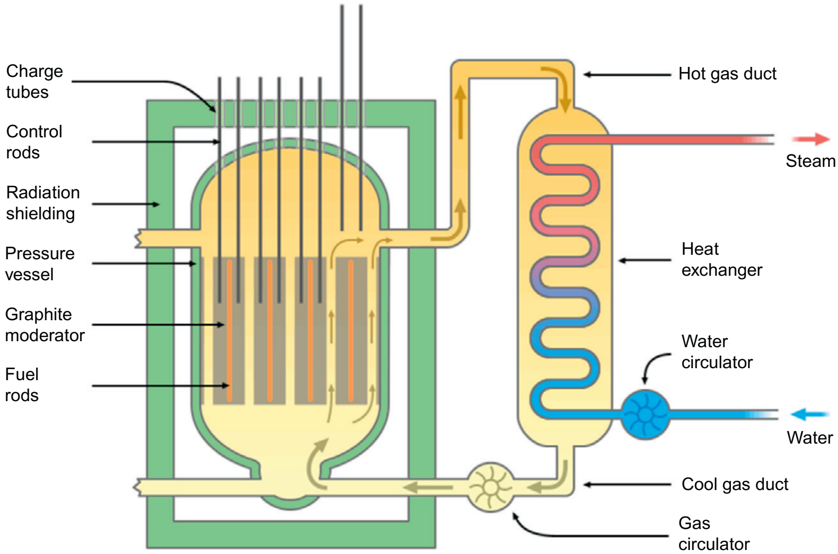

FIGURE 17.4 Schematic of an advanced gas-cooled reactor.

alternative designs. No more U.K. AGR reactors are planned and the last reactor to be built in the United Kingdom was a PWR.

RBMK Reactor

The RBMK reactor is a Russian design that uses a graphite moderator and water as the coolant. Like a BWR, the water is maintained under pressure but allowed to boil around 290 °C and the steam generated is pumped through steam turbines to generate power.

The Soviet Union built 17 of these reactors between 1973 and 1991 during which the design continuously evolved. It was one of these reactors that failed at Chernobyl, with dramatic consequences. It is now known that a major design flaw contributed to the accident. Eleven of these reactors continue to operate, all in Russia.

High-temperature Gas-cooled Reactor

The high-temperature gas-cooled reactor (HTGR) is similar in concept to the AGR (Figure 17.5). It uses uranium fuel, a graphite moderator, and a gas as coolant. In this case, however, the gas is helium.

Several attempts have been made to build reactors of this type but none has so far entered commercial service. Early development work was carried out in the United States. The U.S. design utilized fuel elements in the shape of interlocking hexagonal prisms of graphite containing the fissile material. HTGR fuel

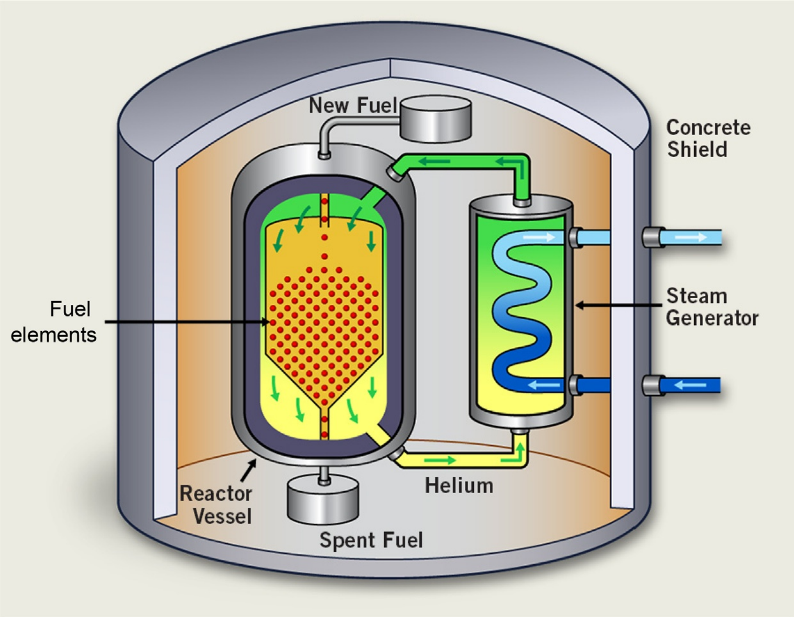

FIGURE 17.5 Schematic of a high-temperature gas-cooled reactor.

is often much more highly enriched than the fuel in a water-cooled reactor, with up to 8% uranium-235. The arrays of hexagonal graphite prisms contain shafts for control rods and passages for the helium to pass through and carry away the heat generated by fission.

Another design, developed in Germany, uses uranium-oxide fuel that is sealed inside a graphite shell to form a billiard ball-sized fuel element called a pebble. This gives the reactor its name: the pebble-bed reactor. Development of this in Germany was eventually abandoned but the idea was taken up during the 1990s by the South African utility Eskom, which continued development until 2010 when it appears to have been halted. Japan and China have funded experimental programs too.

The advantage of the HTGR is that both the moderator, graphite, and the coolant (helium) can operate at high temperatures without reacting or deteriorating. A typical HTGR will operate at a pressure of 100 atm and at a temperature up to 900 °C. This enables better thermodynamic conditions to be achieved, leading to higher efficiency. The reactor is designed so that in the event of a coolant failure it will be able to withstand the rise in internal temperature without failing.

The HTGR can use a dual-cycle system in which the helium coolant passes through a heat exchanger where the heat is transferred to water and steam is generated to drive a steam turbine. This arrangement is around 38% efficient. However, a more advanced system uses the helium directly to drive a gas turbine. This arrangement is sometimes called a gas turbine modular helium reactor (GT-MHR). In theory, the GT-MHR can achieve an energy conversion efficiency of 48%.

One of the advantages of the HTGR is that it can be built in relatively small unit sizes. Modules can have generating capacities between 100 MW and 200 MW, making it attractive for a wider variety of applications. The modular form of most designs also makes it easy to expand a plant by adding new modules. However, no reactors of this design have yet entered commercial service.

Nuclear Fast (Breeder) Reactors

A conventional fission reactor can only use uranium-235 as fuel. Another type of reactor, called a breeder or fast reactor, aims to utilize the much more abundant uranium-238, but since this is stable under normal conditions it has to be approached in a circuitous way (Figure 17.6).

The breeder reactor uses not uranium but plutonium as its primary fuel. This decays or splits when bombarded with neutrons in a similar way to uranium-235, producing fast neutrons and two smaller nuclei. However, whereas in the conventional uranium-235 reactor the fast neutrons must be slowed with

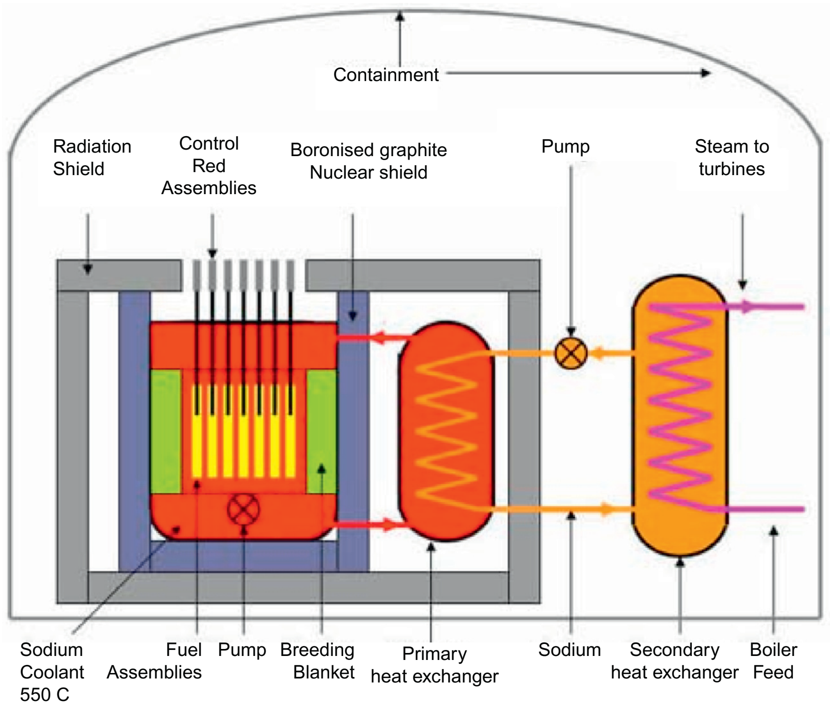

FIGURE 17.6 Schematic of a nuclear fast (breeder) reactor.

a moderator to enable further nuclear fusion reactions to take place, plutonium will interact with the fast neutrons so no moderator is required.

The absence of a moderator means that there is an abundant supply of fast neutrons available. Some of these neutrons can be harnessed to produce more fuel because a uranium-238 atom will interact with a fast neutron to form an atom of plutonium. The core of a breeder reactor is surrounded with uranium-238, and by careful design the reactor can be made to produce more plutonium than it uses, hence the name (the alternative name, fast reactor, comes from its use of fast neutrons).

A breeder reactor needs an initial supply of plutonium. This can be obtained by processing the spent fuel from a conventional reactor where some plutonium is formed when fast neutrons interact with uranium-238 atoms before they are slowed. Once the reactor has started, it should provide its own source of fuel. However, a breeder reactor requires an allied fuel reprocessing plant.

Most breeder reactors use liquid sodium as the coolant because this does not slow the neutrons. A number of prototypes have been built, the most recent in France and Japan. However, the coolant has often proved to be a source of major problems and no commercial breeder reactor has ever been built.

There is a second type of breeder reactor called a thermal breeder reactor. This uses a rare isotope of uranium, uranium-233, as its fissile material. When bombarded with slow neutrons uranium-233 undergoes a fission reaction similar to uranium-235, producing fast neutrons that are then slowed by a moderator. The core of the thermal breeder reactor is surrounded with thorium. This also reacts with slow neutrons to produce more uranium-233. The thermal breeder is simpler than the fast breeder because it can use water both as its moderator to slow the fast neutrons and as its coolant. This design has been developed in India where there are large reserves of thorium.

Advanced Reactors

With the imminent demise of the gas-cooled reactor in the United Kingdom and with no more RBMK reactors likely to be constructed, all the major commercial reactor designs are now water cooled. The PWR, BWR, and CANDU reactors described earlier are all what are known as second-generation designs. Each of these has led to at least one third-generation design and some of these have been or are being built.

The aim of third-generation reactors is to improve the safety of nuclear power generation and at the same time to reduce construction costs. Many of them were originally conceived during the 1980s and available for construction during the 1990s, but none was built until the 21st century.

To make construction simpler, companies have developed standardized designs that they have sought to have certified in different regions of the world. Approval by a nuclear certification authority should mean that, in principle, the planning application for construction of such a plant will be streamlined.

TABLE 17.2 Third-generation Reactors

Reactor	Reactor Type	Generating Capacity (MW)
ABWR	BWR	1371
AP1000	PWR	1117
APR1400	PWR	1450
ESBWR	BWR	1550
EPR	PWR	1650
ACR series	CANDU	700–1200
PBMR	HTGR	180
APWR	PWR	1500–1700
IRIS		100–300

Source: P. Breeze, *The Future of Nuclear Power*, Business Insights, 2007.

Several third-generation reactors have been certified in the Untied States and elsewhere (see Table 17.2).

The main safety advance in third-generation reactors is to design them with passive safety features such that in the event of any type of failure the reactor will be capable of shutting itself down without the need for intervention. Such features are seen as crucial to gaining continued public acceptance of nuclear power.

One of the main new reactors is the advanced boiling water reactor (ABWR), designed by GE and based on its widely used BWR. It has been certified for use in both the United States and in Europe. Four ABWR reactors are operating in Japan and two in Taiwan. Others are planned.

The AP1000 is an evolution of the Westinghouse PWR with passive features intended to reduce construction costs. The first version, the AP600 with a generating capacity of 600 MW, was certified in the United States in 1999. The 1117 MW version, the AP1000, was approved in the United States in 2005. Four units based on this design are under construction in China and the construction of units in the United States has been approved. An evolution of the AP1000 has also been developed for the Chinese market. Called the CAP1400, it has a generating capacity of 1400 MW. China is expected to approve construction of a plant to this design and the Chinese National Nuclear Corp. is also hoping to export the technology.

The APR1400 is based on the Korean Standard Nuclear Power Plant developed by the Korean Electric Power Co. The design of the latter was originally based on the Combustion Engineering System 80+ design, now owned by Westinghouse. Two APR1400 units are under construction in South Korea.

The only third-generation reactor under construction in Europe is the European pressurized water reactor (EPR) also known in the United States as evolutionary pressurized water reactor. This is an evolution of the French PWR together with a German reactor design and has a generating capacity of 1650 MW, the largest of any commercial reactor. The design is being marketed by French company Areva. One is under construction in Finland, one in France, and two in China. The design is seeking certification in the United States and United Kingdom.

The economic simplified BWR (ESBWR) is another evolution of GE's BWR and has been developed by GE and Hitachi. It further refines the passive elements of the ABWR. The design is seeking certification in the United States. None has been built.

The third generation of the CANDU reactor is the ACR series. Two units have been designed: the ACR700 with a generating capacity of 700 MW and the ACR1200 with a capacity of 1200 MW. One of the innovations of the design is to use heavy water in the core but with a light water circuit to extract heat. The reactors would also use uranium enriched to 1.5–2%.

The pebble-bed modular reactor (PBMR) is based on the pebble-bed design described before. It was being developed by a company called PBMR, an affiliate of the South African utility Eskom, but the South African government decided to stop funding for the project in 2010. Its future now looks in doubt.

The advanced PWR (APWR) has been designed by Mitsubishi. It has a potential generating capacity up to 1700 MW. A special version called the US-APWR has been designed for the U.S. market.

A new reactor still under development is the international reactor innovative and secure (IRIS), which is the product of an international consortium involving nine nations. The reactor is small, with a unit size of 100–300 MW. It will have a range of passive features. One of the key innovations is to use uranium enriched to between 5% and 9% so that refueling would only be required every five years.

In addition to these third-generation designs, there are a number of more complex and advanced fourth-generation designs being considered. (Both the PBMR and IRIS might be considered fourth-generation designs.) However, there are economic and political arguments for sticking with the best of the existing second- and third-generation designs since these require minimal fuel recycling and provide less danger of nuclear proliferation.

NUCLEAR FUSION

Nuclear fusion, the reaction that fuels the sun and the stars, has excited scientists and technologists ever since the process was identified during the 1930s. Unsuccessful attempts at fusion took place during the 1930s but halted during World War II. Experimental work restarted during the late 1940s. Since then a series of fusion reactors have been built around the world. Around 20 are in operation today.

In 1958 at an Atoms for Peace conference in Geneva, fusion research was established as an international collaborative venture and at least one strand of fusion development, which based on magnetic confinement, has remained international in flavor ever since. The necessity for this was reinforced during the 1970s when it became clear that the cost of developing fusion was likely to be beyond the resources of any one nation.

While large fusion reactors based on magnetic confinement were being built in the United States, Europe, and Japan, other developments remained hidden behind the security of nuclear armaments research. Fusion is the basis for the hydrogen bomb and so much of the research into its development and control remained secret until very recently. It is this research that has led to the idea of inertial confinement, an entirely different approach to fusion for power generation. During the last five years the veil of secrecy has at least partly dropped and a major program in the United States aims to develop a demonstration power plant during the 2020s. Meanwhile, the largest international magnetic confinement reactor is under construction in the south of France and should start experiments at around the same time.

Magnetic Confinement

The fusion reaction between deuterium and tritium (DT) discussed earlier is the easiest to achieve in a reactor, but even so it requires extremes of both temperature and pressure. If it can be mastered, then potentially fusion could provide almost limitless amounts of energy. In theory, 1 tonne of deuterium could provide the equivalent of 3×10^{10} tonnes of coal.

The temperature required to achieve fusion with DT is over 100×10^6 °C. Under these conditions the atoms disintegrate to create a sea of electrons and nuclei, a fourth state of matter called a plasma. There are no materials in existence that can survive the plasma temperature, so an alternative way has to be found to contain and control the plasma. The most promising solution is by means of a magnetic field (Figure 17.7).

Magnetic containment was recognized early in fusion research as the only way to maintain a fusion plasma, but it was not until the 1950s that the best form of magnetic field, the torroidal field, was identified by scientists in Russia. Here a device called a tokamak was developed, and in the early 1960s experimental results showed that the high temperatures required for fusion could be achieved with this device.

Since then a series of ever larger tokamak reactors have been constructed. The most important of these were the joint European torus (JET) at Culham, UK, Tokamak fusion test reactor (TFTR) at Princeton, NJ, JT-60U in Naka, Japan, and T-15 in Moscow, Russia. Both TFTR and JET experimented with

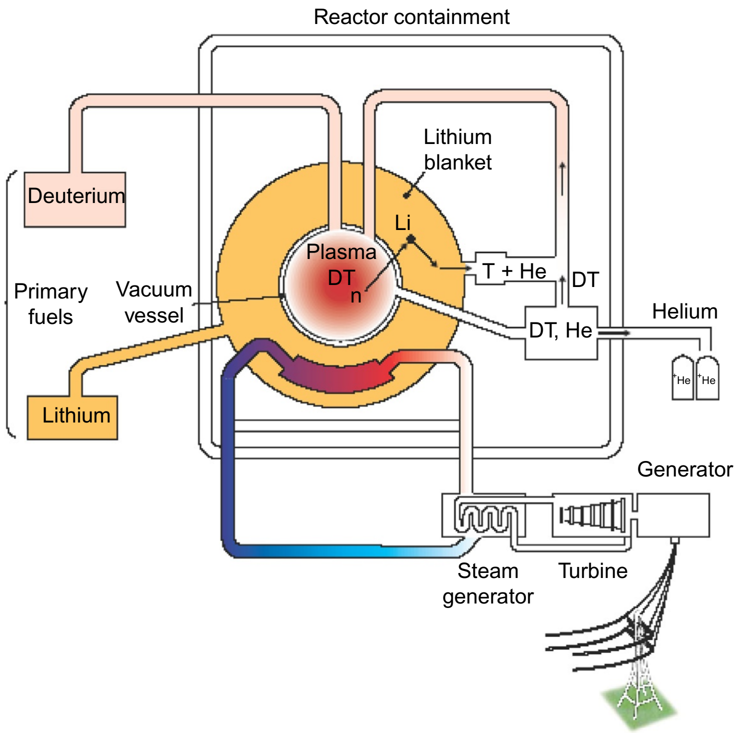

FIGURE 17.7 Schematic of a fusion power plant based on a magnetic confinement reactor.

DT fuel from the beginning of the 1990s, and in 1997 JET established the record for the greatest amount of energy generated by a fusion reactor, 16 MW.

Even with this output, the reactor consumed more energy than it generated. JET achieved a power-in-to-power-out ratio (the gain of the reactor) of around 0.7. A gain of 1 represents the breakeven point. Even more crucially, JET could only maintain the plasma burst for five seconds. If it continued for longer, its systems would begin to overheat.

The next stage in magnetic confinement fusion development is the international thermonuclear experimental reactor (ITER—the word is also Latin for "the way"). This project was first conceived in 1988 and an agreement to build the plant was finally signed in 2007 by the EU, Russia, Japan, United States, China, India, and South Korea.

ITER will have a plasma volume of 800 m^3 and a power output of 500 MW_{th}, 30 times that of JET. Fusion energy generation is a matter of size and at this size it is hoped that ITER will have a gain of 10, producing 500 MW_{th} from an input of 50 MW_{th}. This will be sufficient to prove fusion as a net source of energy, but ITER has not been designed to generate power, so it will not have all the features needed for a demonstration plant. That will have to wait until the successor to ITER.

Inertial Confinement

While magnetic confinement seeks to create a stable continuous plasma in which fusion can take place, the alternative—inertial confinement—seeks instead to generate energy from a series of discrete fusion reactions producing a burst of energy each time (Figure 17.8). In an inertial confinement reactor, small capsules containing around 150 mg of a mixture of deuterium and tritium (DT) are exposed to a massive pulse of energy from multiple lasers. When the laser beams strike the capsule they create an explosion of X-rays from its surface, and these in turn (by the mechanical principle of action and reaction) create a pressure pulse that heats and compresses the DT mixture with such vigor that the conditions for fusion are generated at its core. One fusion starts, the reaction radiates outwards through the DT mixture faster than the actual molecules can expand and escape (they are "confined" by their inertia), and so the whole charge undergoes fusion and releases a pulse of energy.

To make this into a means of generating power, these small exploding suns must be created at a relatively rapid rate of perhaps 15 each second. This sounds both exacting and ambitious, but it is exactly what a U.S. program proposes. To achieve it, the U.S. government has built the National Ignition Facility (NIF), a $5 billion project that is intended to serve both military and civilian research.

NIF is provided with 192 lasers capable of providing a pulse of up to 5 MJ of energy and it has so far produced 1.8 GJ, equivalent to a delivery rate of 500 TW of energy. All the energy contained in the laser beams is focused onto the 150 mg of DT. NIF is only capable of single-shot experiments rather than the continuous operation required for a power plant. Since it started in 2009 it has carried out a series of experiments trying to achieve ignition, the point at which the DT produces more fusion energy than the lasers pump into it. By 2013 it was a factor of two or three short of ignition.

Alongside NIF is the Laser Inertial Fusion Energy (LIFE) project, a collaboration of scientists, technologists, utilities, and regulators that are seeking to design a power plant capable of exploiting inertial confinement. Current plants see a demonstration project constructed between 2020 and 2030 and commercial plants available by 2030 or soon afterwards.

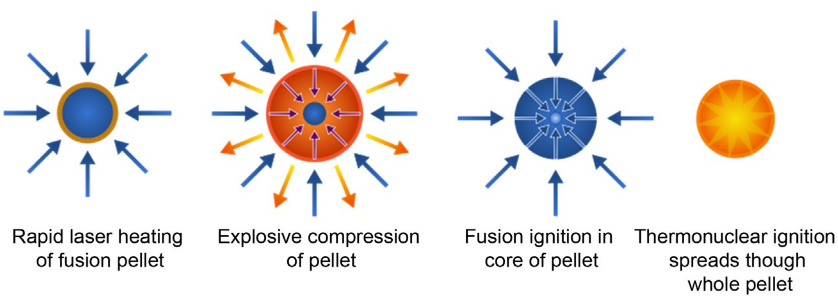

FIGURE 17.8 Principle of inertial confinement.

Tritium Production

For either magnetic confinement or inertial confinement to become successful sources of electrical power, each needs a source of tritium. This can be generated in a nuclear reaction from lithium. To provide self-sustaining power plants, each type of plant will have to produce its own tritium. Under current designs it is conceived that this will be manufactured in a "blanket" that surrounds the core of the fusion reactor, much in the same way as breeder reactors produce their own fuel. Liquid lithium might also be used as the reactor coolant although alternatives might prove easier to manage.

Designing the blanket and heat-recovery system is one of a number of major hurdles that have to be jumped if fusion is to become viable commercially. While the goal is still distant, it seems more likely to be achieved than it has at any stage in the past.

ENVIRONMENTAL QUESTION

The use of nuclear power raises important environmental questions. It is an apparent failure to tackle these satisfactorily that has led to much of the popular disapprobation that the nuclear industry attracts. There are two adjuncts to nuclear generation that cause the greatest concern: nuclear weapons and nuclear waste.

While the nuclear industry would claim that the civilian use of nuclear power is a separate issue to that of atomic weapons, the situation is not that clear-cut. Nuclear reactors are the source of the plutonium that is a primary constituent of modern nuclear weapons. Plutonium creation depends on the reactor design; a breeder reactor can produce large quantities while a PWR produces very little. Nevertheless, all reactors produce waste that contains dangerous fissile material. This is a subject of international concern.

The danger is widely recognized. Part of the role of the International Atomic Energy Agency is to monitor nuclear reactors and track their inventories of nuclear material to ensure than none is being sidetracked into nuclear weapons construction. Unfortunately, this system can never be foolproof. It seems that only if all nations can be persuaded to abandon nuclear weapons can this danger, or at least the popular fear of it, be removed. Such an agreement looks highly improbable.

The problem is political in nature. Nevertheless, it carries a stigma from which the industry can never escape. The prospect of a nuclear war terrifies most people. Unfortunately, for the nuclear power industry, some of the after-effects of nuclear explosion can also be produced by a major civilian nuclear accident. The contents of a nuclear reactor core include significant quantities of extremely radioactive nuclei. If these were released during a nuclear accident they would almost inevitably find their way into humans and animals via the atmosphere or through the food chain.

Large doses of radioactivity or exposure to large quantities of radioactive material kills relatively swiftly. Smaller quantities of radioactive material are lethal too, but over longer time scales. The most insidious effect is the genesis of a wide variety of cancers, many of which may not become apparent for 20 years or more. Other effects include genetic mutation that can lead to birth defects.

The prospect of an accident leading to a major release of radio-nuclides has created a great deal of apprehension about nuclear power. The industry has gone to extreme lengths to tackle this apprehension by building ever-more sophisticated safety features into their power plants. Unfortunately, the accidents at Three Mile Island in the United States, Chernobyl in the Ukraine, and Fukushima Daiichi in Japan remain potent symbols of the danger.

This danger has been magnified by a rise in international terrorism. The threat now exists that a terrorist organization might cause a nuclear power plant accident or, by exploiting contraband radioactive waste or fissile material, cause widespread nuclear contamination.

So far a nuclear incident of catastrophic proportions has been avoided, though both Chernobyl and Fukushima have caused extensive disruption, and in the case of the former, a disputed number of deaths as a result of radioactive exposure. Smaller incidents have been more common and low-level releases of radioactive material have taken place. The effects of low levels of radioactivity have proved difficult to quantify. Safe exposure levels are used by industry and regulators but these have been widely disputed. Only the elimination of radioactive releases from civilian power stations is likely to satisfy a large sector of the public.

Radioactive Waste

As the uranium fuel within a nuclear reactor undergoes fission, it generates a cocktail of radioactive atoms within the fuel pellets. Eventually the fissile uranium becomes of too low a concentration to sustain a nuclear reaction. At this point the fuel rod will be removed from the reactor. It must now be disposed of in a safe manner. Yet after more than 60 years of nuclear fission, no safe method of disposal has been developed.

Radioactive waste disposal has become one of the key environmental battlegrounds over which the future of nuclear power has been fought. Environmentalists argue that no system of waste disposal can be absolutely safe, either now nor in the future. And since some radio-nuclides will remain a danger for thousands of years, the future is an important consideration.

Governments and the nuclear industry have tried to find acceptable solutions. But in countries where popular opinion is taken into consideration, no mutually acceptable solution has been found. As a result, most spent fuel has been stored in the nuclear power plants where it was produced. This is now causing its own problems as storage ponds designed to store a few years' waste become filled, or overflowing.

One avenue that has been explored is the reprocessing of spent fuel to remove the active ingredients. Some of the recovered material can be recycled as fuel. The remainder must be stored safely until it has become inactive. But reprocessing has proved expensive and can exacerbate the problem of disposal rather than assisting it. As a result, it too appears publicly unacceptable.

The primary alternative is to bury waste deep underground in a manner that will prevent it ever being released. This requires both a means to encapsulate the waste and a place to store the waste once encapsulated. Encapsulation techniques include sealing the waste in a glasslike matrix. Finding a site for such encapsulated waste has proved problematic. An underground site must be in stable rock formation in a region not subject to seismic disturbance. Sites in the United States and Europe have been studied but none has yet been accepted. Even if site approval is achieved, there appears little prospect of any nuclear waste repository being built until well into the middle of the 21st century.

Other solutions have been proposed for nuclear waste disposal. One involves loading the fuel into a rocket and shooting it into the sun. Another utilizes particle accelerators to destroy the radioactive material generated during fission.

Environmentalists argue that the problem of nuclear waste is insoluble and represents an ever-growing burden on future generations. The industry counters this, but in the absence of a persuasive solution its arguments lack weight. Unless a solution is found, the industry will continue to suffer.

Waste Categories

Spent nuclear fuel and the waste from reprocessing plants represent the most dangerous of radioactive wastes, but there are other types too. In the United States these first two types of waste are categorized as high-level waste[4] while the remainder of the waste from nuclear power plant operations is classified as low-level waste. There is also a category called transuranic waste, which is waste containing traces of elements with atomic numbers greater than that of uranium (92). Low-level wastes are further subdivided into classes depending on the amount of radioactivity per unit volume they contain.

In the United Kingdom there are three categories of waste: high level, intermediate level, and low level. High level includes spent fuel and reprocessing plant waste, intermediate level is mainly the metal cases from fuel rods, and low level constitutes the remainder. Normally both high- and intermediate-level waste require some form of screening to protect workers, while low-level waste can be handled without a protective radioactive screen.

4. The U.S. Department of Energy does not classify spent fuel as waste but the Nuclear Regulatory Commission does.

High-level wastes are expected to remain radioactive for thousands of years. It is these wastes that cause the greatest concern and for which some storage or disposal solution is most urgently required. But these wastes form a very small part of the nuclear waste generated by the industry. Most is low-level waste. Even so, it too must be disposed of safely. Low-level waste can arise from many sources. Anything within a nuclear power plant that has even the smallest exposer to any radioactive material must be considered contaminated. One of the greatest sources of such waste is the fabric of a nuclear power plant itself.

Decommissioning

A nuclear power plant will eventually reach the end of its life, and when it does, it must be decommissioned. At this stage the final, and perhaps largest, nuclear waste problem arises. After 30 or more years[5] of generating power from nuclear fission, most of the components of the plant have become contaminated and must be treated as radioactive waste. This presents a problem that is enormous in scale and costly in both manpower and financial terms.

The cleanest solution is to completely dismantle the plant and dispose of the radioactive debris safely. This is also the most expensive option. A halfway solution is to remove the most radioactive components and then seal up the plant for 20–50 years, allowing the low-level waste to decay, before tackling the rest. Two Magnox reactor buildings in the United Kingdom were sealed in this way in 2011 and are expected to remain in that state for 65 years. A third solution is to seal the plant up with everything inside and leave it, entombed, for hundreds of years. This has been the fate of the Chernobyl plant.

Decommissioning is a costly process. Regulations in many countries now require that a nuclear-generating company put by sufficient funds to pay for decommissioning of its plants. The U.S. utility Southern California Edison has put aside $2.7 billion to decommission its San Onofre power plant, expecting this to cover around 90% of the total expenditure. Meanwhile, in 2011 the U.K. government estimated nuclear decommissioning costs for its existing power plants to be £54 billion. When building a new nuclear plant, the cost of decommissioning must, therefore, be taken into account.

COST OF NUCLEAR POWER

Nuclear power is capital intensive and costs have escalated since the early days of its development. This is partly a result of higher material costs and high interest rates, but it is also a result of the need to use specialized construction materials and techniques to ensure plant safety. In the United States, in the early

5. U.S. nuclear plants are now winning operating license extensions to allow them to continue operations for up to 60 years.

1970s, nuclear plants were being built for units costs of $150–300/kW. By the late 1980s, the figures were $1000–3000/kW.

The Taiwan Power Company carried out a study, published in 1991, which examined the cost of building a fourth nuclear power plant in Taiwan. The study found that the cost for the two-unit plant would be U.S. $6.3 billion, a unit cost of around $3150/kWh. The estimate was based on completion dates of 2001 and 2002 for the two units. Orders were actually placed in 1996, with construction scheduled for completion in 2004 and 2005. Construction actually started in 1999 and the plants were still not completed in 2013.

In its 2012 Annual Energy Outlook, the U.S. Energy Information Administration (EIA) estimated that the overnight cost of an advanced nuclear plant based on an order in 2011 for a plant that would enter into service in 2017 was $4619/kW. When contingency factors were taken into account, this rose to $5335/kW.[6] Meanwhile, a 2011 U.K. study put the capital cost of nuclear power in the United Kingdom at £3500/kW.[7] Nuclear construction costs do not always take into account decommissioning. This can cost from 9% to 15% of the initial capital cost of the plant.

The fuel costs for nuclear power are much lower than for fossil fuel–fired plants, even when the cost of reprocessing or disposal of the spent fuel are taken into account. Thus, levelized costs of electricity provide a more meaningful picture of the economics of nuclear power generation.

In the 2012 Annual Energy Outlook the EIA estimated that the cost of electricity from a new nuclear power plant entering service in 2017 would be $113/MWh. This was similar to the EIA's estimated cost of electricity from an advanced coal-fired power plant (without carbon capture and storage) but more expensive than gas-fired combined cycle generation, even with carbon capture and storage. From the U.K. 2011 study the cost of electricity from a nuclear plant was around £97/MWh, cheaper than either coal- or natural gas–fired plants with carbon capture and storage.

While the cost of new nuclear-generating capacity might be considered expensive in some parts of the world, but acceptable in others, the cost of power from existing nuclear power plants is often extremely competitive. This is true even where coal and gas are readily available. In support of this, a number of companies have, in the 21st century, started to make a successful business of running U.S. nuclear power stations sold by utilities when the U.S. industry was deregulated. In France, too, nuclear power is on average the cheapest source of electricity. Here, however, it may be considered a nationalized industry.

---

6. Assumptions to the Annual Energy Outlook 2012, U.S. Energy Information Administration.
7. M. MacDonald, *Costs of Low Carbon Generation Technologies*, U.K. Committee on Climate Change, 2011.

Index

Note: Page numbers followed by *f* indicate figures, *t* indicate tables and *np* indicate footnote.

A

B

F

G

H

I

N

R

U

V

W

Y

Z