Chapter 9
Engineering: The Big Challenge*

Introduction

With the information they have gathered from the public media, most people who have heard of fusion consider fusion energy to be a pipedream. Their information is out of date. As we have shown in the last two chapters, great advances have been made in fusion physics, and our knowledge of plasma behavior in toroidal magnetic bottles is good enough for us to push on to the next step. This does not mean, however, that fusion is *not* a pipedream. There is a large chasm between the understanding of the physics and the engineering of a working reactor. There are problems in the technology of fusion so serious that we do not know if they can be solved. But the payoff is so great that we have to try.

The situation can be compared – or contrasted – with that of the Apollo program to put a man on the moon. In that program, the physics was already known: Newton's laws of motion covered all the physics that was needed. In the case of fusion, it took over 50 years to establish the science of plasma physics, to develop fast computers, and to understand the physics of magnetic confinement; but we have done it. In the Apollo case, there were engineering problems whose solutions could not be fully tested. Could the nose cone material stand up to the heat of reentry? Can humans survive long periods without gravity and then the stress of reentry? Will micrometeorites puncture the space suits of the astronauts? It was a dangerous experiment, but President Kennedy pushed ahead, and it succeeded marvelously. In the case of fusion, we do not know yet how to build each part of a reactor, but the only way to get this ideal source of energy is to push on ahead. The expense will be comparable to Apollo's, but at least no human lives are endangered.

The path to a commercial fusion reactor has been studied intensely in the past decade. There are three or four steps: (1) the ITER experiment now being built, (2) one or more large machines for solving engineering problems, (3) DEMO, a prototype reactor built to run like a real reactor but not producing full power, and (4) FPP, fusion power plant, a full-size reactor built and operated by the utilities industry.

---

*Numbers in superscripts indicate Notes and square brackets [] indicate References at the end of this chapter.

F.F. Chen, *An Indispensable Truth: How Fusion Power Can Save the Planet*,
DOI 10.1007/978-1-4419-7820-2_9, © Springer Science+Business Media, LLC 2011

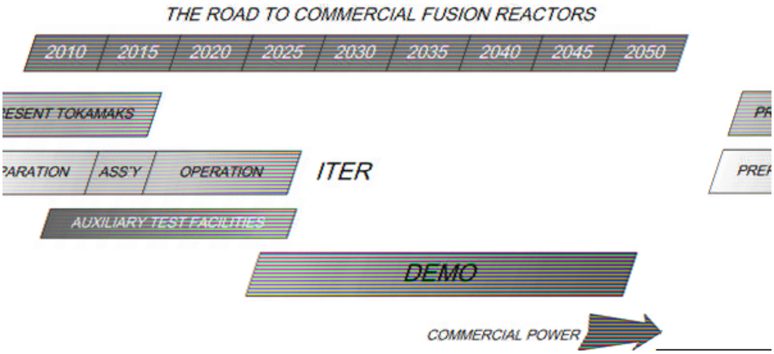

Fig. 9.1 A possible schedule for developing fusion power (Data from G. Janeschitz, *The physics and technology basis of ITER and its mission on the path to DEMO*, Symposium on Fusion Energy, San Diego, California, June 2009)

Step 2 is being hotly debated. Some think that experiments on ITER will give enough information to design DEMO. Others propose intermediate machines designed to solve specific problems such as the tokamak wall material or the breeding of tritium. These problems are described in the main part of this chapter. The time it will take to reach the FPP stage might look something like this (Fig. 9.1). Any additional machines for engineering testing before designing DEMO are shown in Fig. 9.1, although they may not be necessary. Although this timeline is called the "fast track" to fusion, it still will take until 2050 before fusion power becomes a reality. The economic downturn at the turn of this decade has already delayed the construction of ITER. Shortening this timeline can be done only with greatly increased funding. In the meantime, expansion of the other renewable energy sources listed in Chap. 3 is still necessary.

The two toughest engineering problems are the material of the "first wall" and the breeding of tritium. These will be discussed in detail. We also mentioned some physics problems that are not completely solved. One concerns "disruptions" which kill the plasma and must be avoided in a reactor. The best known way to avoid them is to operate safely below the tokamak's limits, and this means less output power. Otherwise, injection of a large puff of gas can stop an incipient disruption; this is a crude solution. A second problem concerns the edge-localized modes (ELMs), instabilities that dump plasma energy into places not designed to absorb it. Currently, internal correction coils are to be inserted inside the plasma chamber to suppress ELMs as well as resistive wall modes (RWMs). This is another crude solution which would not be suitable in a reactor. A third problem concerns the alpha particles (the helium nuclei) which are the products of the D–T fusion reaction. These fast ions can, in theory, excite Alfvén waves, and these electromagnetic waves could disrupt plasma confinement. This instability cannot be studied until we can ignite a plasma to produce these alpha particles.

Although these seem to be formidable problems, there will be a learning curve when ITER and DEMO are built. Once industry gets serious about fusion, progress will be rapid. We will go from Model-T Fords to Mercedes-Benzes. We will go from

DC-3s to Airbus A380s. We may even get lucky with more help from Mother Nature and find that fast alpha particles are stabilizing. Where there's a will, there's a way. With a positive attitude, the fusion community can continue to achieve and live up to its track record of the last 50 years. Further in the future, in the second half of this century, a second generation of fusion reactors will look quite different from the tokamak as described here. There are other magnetic configurations, simpler than the tokamak, that have not been fully developed for lack of funding. These are described in Chap. 10. Better yet, there are fuel cycles that do not require tritium, thus avoiding almost all of the fuel breeding and radioactivity problems of the first generation of fusion reactors. These advanced fuel cycles can run only with hotter and denser plasmas than we can now produce, but which may be possible once we have learned how to control plasma better. Advanced fuels are also presented in Chap. 10. The engineering problems described here are not the end of the story.

The First Wall and Other Materials

The First Wall

Figure 9.2 is a more realistic drawing of the ITER machine than shown in Chap. 8. The plasma will occupy the D-shaped vacuum space surrounded by tiles. These tiles are the plasma facing components (PFCs), commonly called the "first wall." They have to withstand a tremendous amount of heat from the plasma and yet must not contaminate the plasma and be compatible with the fusion products that impinge on them. Early tokamaks used stainless steel, but clearly this is not a high-temperature material. Current tokamaks use carbon fiber composites (CFCs), a light, strong, high-temperature material that is used in bicycles, racing cars, and space shuttles. Just as rebars are used to strengthen concrete, carbon fibers are used to strengthen graphite. However, carbon cannot be used in a reactor because it absorbs tritium, which would not only deplete this scarce fuel but also weaken the CFC. After all, hydrocarbons like methane and propane are very common, stable compounds; and tritium is just another form of hydrogen and can be captured by the carbon to form hydrocarbons.

Tungsten is a refractory metal, but it is high-Z; that is, it has a high atomic number and therefore has so many electrons that it cannot be completely ionized. The remaining electrons radiate energy away, cooling the plasma. The good thing about hydrogen and its isotopes is that they have only one electron, and once that electron is stripped free of the nucleus by ionization, the atom can no longer emit light. Beryllium is a suitable low-Z material, but it has a low melting point, and so has to be cooled aggressively. In preparation for ITER, the European tokamak JET is being upgraded with a beryllium first wall. In short, the first-wall material must not absorb tritium and must have a low atomic number, take high temperatures, and be resistant to erosion, sputtering, and neutron damage.

ITER, of course, is only the first step. There are large steps between ITER and DEMO and between DEMO and a full reactor. Some large numbers on the first wall

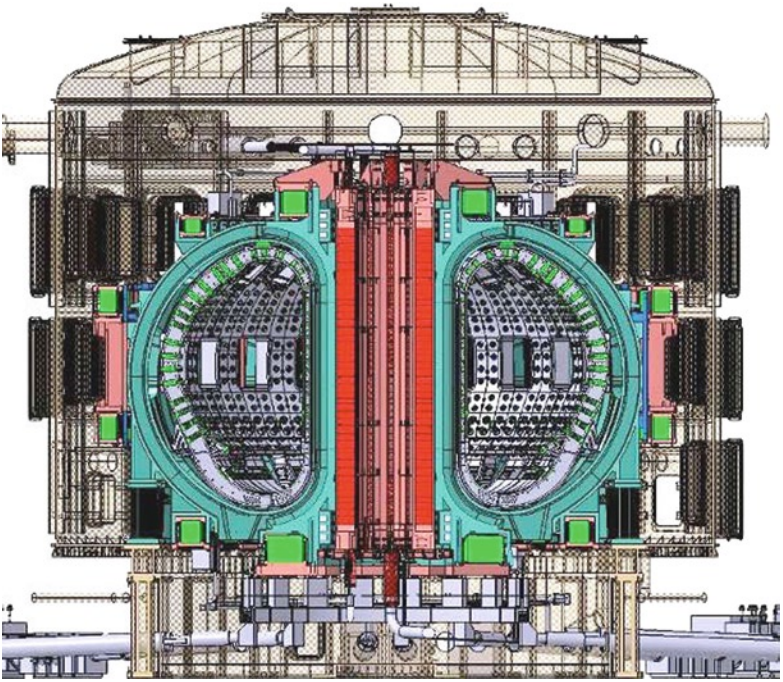

Fig. 9.2 Diagram of ITER, showing the "first wall" and openings (ports) where experimental modules can be inserted for testing [29]

Table 9.1 Loads on the first wall

	ITER	DEMO	Reactor	Units
Fusion power	0.5	2.5	5	GW
Heat flux	0.3	0.5	0.5	MW/m^2
Neutron load	0.78	<2	2	MW/m^2
Neutron load in life	0.07	8	15	MW-years/m^2
Neutron damage	<3	80	150	dpa

are given in Table 9.1. We see that the step between ITER and DEMO is much larger than between DEMO and Reactor. Hence the call for a materials-testing facility intermediate between ITER and DEMO.

The fusion power is given in gigawatts. A typical power plant generates 1 GW of electricity; and perhaps 5 GW of fusion power is needed to give that, since the tokamak needs power to run, and there is still a heat cycle in a steam plant to produce electricity. The heat flux impinging on the first wall is about 0.5 MW/m^2. This translates to 50 W/cm^2 or about 300 W/sq. in. This is not much more than the surface of an electric iron, though the total heat is considerable. The real problem is in the divertor, which has to handle most of the heat from the plasma. Divertors will be covered later.

The neutron load is the energy of the 14-MeV neutrons from the D–T reaction which pass through the first wall. This energy is not deposited in the first wall, but the neutrons damage the wall. The neutron load summed over the life of the wall is what matters. This is much larger for a reactor than for ITER, since ITER is just an experiment, while a reactor should last about 15 years before it has to be revamped. The neutron damage is measured in displacements per atom (dpa). The longer the material is exposed to a neutron flux, the more times one of its atoms will be knocked out of place by a neutron. After many dpa's, the material will swell or shrink and become so brittles as to be useless.

Beryllium melts so easily that it cannot be used in a reactor. Boron coating has been tried successfully, but also cannot take high temperatures. Tungsten seems to be the best available wall material because it does not erode or sputter easily and has a high melting point of 3,410°C. However, it is a high-Z material and also cannot be machined easily. A liquid lithium first wall has been considered, but it is no longer proposed.[1] Silicon carbide (SiC) is a promising material that has been studied extensively in the laboratory but does not have a known method for manufacturing in large quantities [1]. How SiC compares with other materials in operating temperature is shown in Fig. 9.3. These temperature ranges are for irradiated materials so that the swelling and fracture effects caused by neutrons are included. Carbon fiber-reinforced graphite (C/C) can take high temperatures, but carbon cannot be used because of tritium retention. Tungsten and molybdenum are classical refractory metals but will cool the plasma if they sputter into it. Silicon carbide reinforced with layers of SiC fibers (SiC/SiC) seems to be the ideal material for the first wall if it can be made without impurities. It takes high temperature, is quite strong, and is resistant to radiation damage. It can last for the 15-year life of the

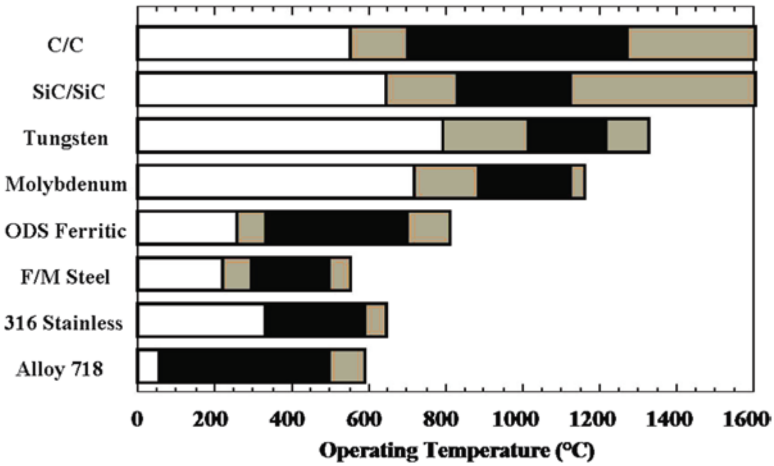

Fig. 9.3 Temperature range of various wall materials under irradiation [1]. The *top* four are refractory materials, and the *bottom* four are structural steels. The *dark center* of each bar is a reasonable operating range; the total bar is an extended range which is possible but not proven

reactor. Its properties have been measured in fission reactors [2]. The main drawback is a thermal conductivity lower than for other materials.

The latest high-tech material is a SiC matrix/graphite fiber composite [1], which has increased thermal conductivity in addition to the other good properties. These advanced materials cannot be designed with existing computer programs, which are applied only to metals. Some reactor studies assume that SiC first-wall material will be available. Though SiC composites have tremendous potential, much research and testing remain to be done before they become a reality.

The Divertor

Sixty percent of the plasma exhaust is designed to go into the "divertor," thus sparing the first wall from the major part of the heat load. Materials and cooling methods can be used in the divertor that cannot be used for the first wall. Figure 9.4 shows how this is done. Special coils located at the bottom of the chamber bend the outermost field lines so that they leave the main volume and enter the divertor. Plasma tends to follow the field lines, so that most of it leaves the chamber by striking the surfaces of the divertor rather than the first wall. Only

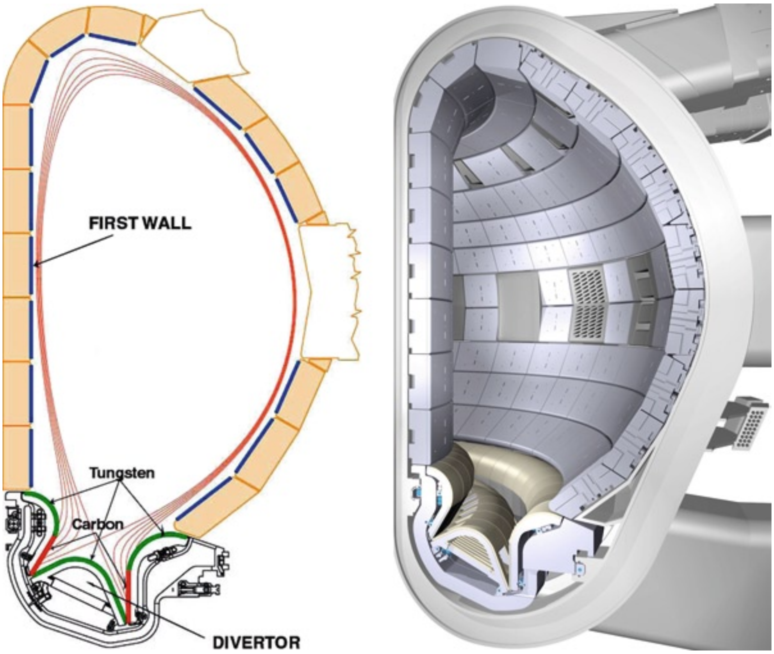

Fig. 9.4 Two views of a tokamak cross section showing the divertor, the first wall, and some ports for heating and diagnostics equipment or for test modules [30, 31]. In the *left diagram*, the outermost magnetic field lines are drawn, showing how they lead the plasma into the divertor. The closed magnetic surfaces in the interior have been omitted for clarity

the plasma that migrates *across* the magnetic field hits the first wall. The heat load on the first wall can be larger than average when there is an instability such as an ELM or a disruption that takes plasma across the field lines suddenly. The first wall of ITER will have to withstand such heat pulses, but DEMO must be built to avoid such catastrophes.

As can be seen in the diagram, the boundary layer of diverted field lines is very thin, only about 6 cm in ITER. In the divertor, these field lines are spread out over a larger area, and the surfaces which the plasma strikes are inclined almost parallel to the field lines so that the heat is deposited over as large a surface as possible. Tungsten can be used for these surfaces, and even carbon compounds can be used in spite of their tritium retention. The divertor parts are easier to replace than the first wall, so the tritium can be removed periodically. The heat load on the divertor surfaces is huge, some 20 MW/m^2, so the cooling system is an important part of the design. Water cooling is possible in ITER, but helium cooling at higher temperatures would have to be used in DEMO and FPPs. The conditions inside a divertor are so intense that they are hard to imagine. Ions with tens of keV energy stream in along the field lines, accompanied by electrons that neutralize their charges. When the ions meet a solid surface, they recombine with electrons to form neutral atoms. There is a dense mixture of plasma with neutral gas made of deuterium, tritium, helium, and impurities, which later have to be separated out in an exhaust processing unit. The neutral gas has to be pumped away fast by vacuum pumps before it flows back into the main chamber and gets ionized again into ions and electrons. To trap the neutrals inside the divertor, a dome-shaped structure has to be added. Figure 9.5 shows the main parts of a divertor designed for ITER. The plasma impinges at a glancing angle onto the high-temperature surfaces made of tungsten and CFC. A heat-sink material, CuCrZr, transfers the heat to water-cooled surfaces.

Water cooling, which is limited to about 170°C, would be insufficient for DEMO and FPP, and cooling by helium gas would have to be used. The helium

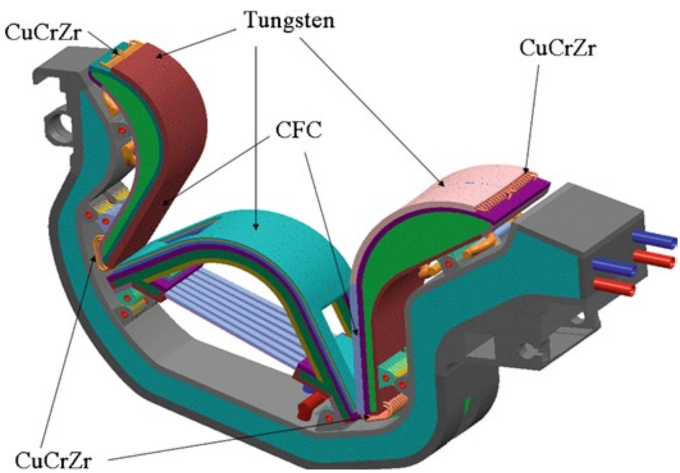

Fig. 9.5 Conceptual diagram of a water-cooled divertor [31]

would be injected at 540°C and be heated to 720°C, while the tungsten and CFC tiles would get to 2,500°C [3]. The coolant would be injected under pressure to cool a small dome as illustrated in Fig. 9.6. These domes are then assembled into nine-finger units, and these units then form a uniformly cooled surface.

Divertor technology is in better shape than other problem areas because divertors are small, and they have already been extensively tested. For instance, meter-sized tungsten and CFC divertor segments (Fig. 9.7) have been tested in Karlsruhe, Germany, up to heat fluxes of 20 MW/m^2. In that large laboratory, divertor materials

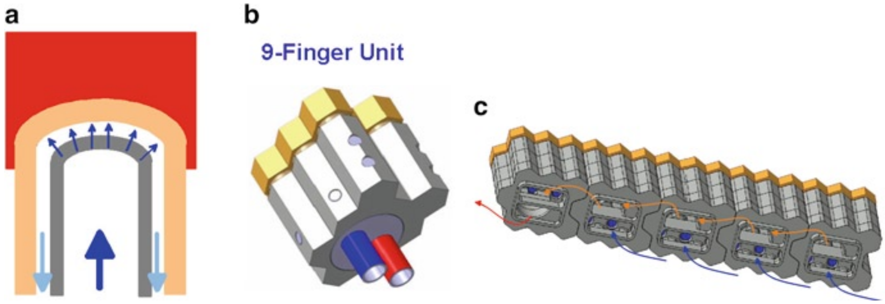

Fig. 9.6 Possible design of a helium cooling system for a divertor [31]. Helium cools a dome-shaped "finger" (**a**), and nine of these are assembled into one unit (**b**). A number of these together then form a cooled surface (**c**)

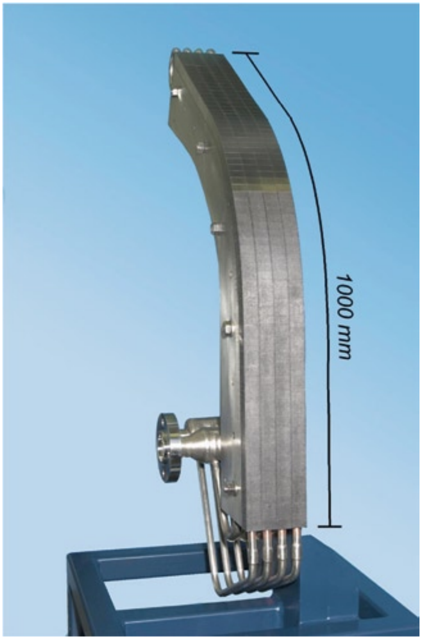

Fig. 9.7 A water-cooled divertor test surface [31]

have been neutron-irradiated, and their manufacturing and assembly techniques have been worked out. Even remote handing techniques for replacing divertors have been tested. It seems possible to design water-cooled divertors for heat fluxes up to 20 MW/m^2 and helium-cooled divertors up to 15 MW/m^2 [31].

Structural Materials

Aside from materials exposed to plasma and large heat fluxes, structural materials have to be chosen to support the huge weight of the reactor elements – the vacuum chamber, magnetic coils, breeding blankets, and so forth. Normally one would use steel; but for fusion, the type of steel has to be carefully designed. The neutrons bombarding the structure will make it radioactive. Only the following elements can be used: iron, vanadium, chromium, yttrium, silicon, carbon, tantalum, and tungsten. Elements like manganese, titanium, and niobium used in other steels would result in long-lived radioactive isotopes. Two *Reduced Activation Ferritic Martensitic Steels* have been designed: Eurofer (in Europe) and F82H (in Japan). These have the following additives to iron [4]:

	Chromium (%)	Tungsten (%)	Vanadium (%)	Tantalum (%)	Carbon (%)
Eurofer	7.7	2	0.2	0.04	0.09
F82H	8.9	1	0.2	0.14	0.12

These steels have only short-lived radioactivity and, unlike fission products, are nonvolatile and can be re-used after storage for 50–100 years. The amount of swelling under neutron bombardment is much smaller than for ordinary stainless steel. Swelling and embrittlement come from helium and hydrogen bubbles trapped in the steel. There are experimental oxide dispersion strengthened (ODS) steels which have nanoparticles of Y_2O_3 that can trap helium and hydrogen, strengthen the material, and reduce creep. Though much has to be done to manufacture these materials with low impurity levels, to study their welding properties, and to test their limits in temperature and radiation resistance in full-time operation, structural materials are not one of the worrisome problems in fusion technology.

Figure 9.8 shows the predicted radioactivity of Eurofer and SiC in a fusion reactor after 25 years of full-power operation. Note that the scales are logarithmic, so that each vertical division represents a factor of 10, and each horizontal division a factor of 100. After 100 years, the radioactivity has decayed by a factor of almost 1,000,000. This material is solid and will not leak out of its containers. The main danger from radioactivity comes from tritium, which decays in 12 years and will be considered in detail later. Note that even this small amount of radioactivity compared with fission is caused by the fact that the D–T reaction emits energetic neutrons. In second-generation fusion reactors using advanced fuels there will be almost no radioactivity.

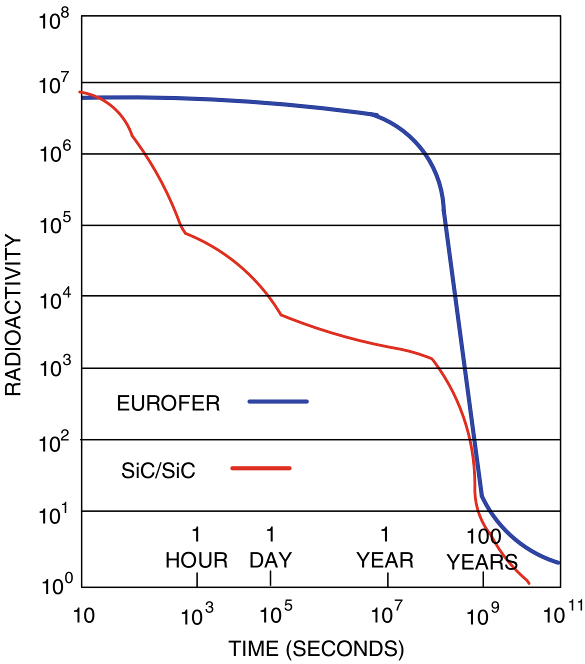

Fig. 9.8 Predicted radioactivity of Eurofer steel and silicon carbide composites versus time. The vertical axis is in units of Becquerels per kilogram. Data from [4]

Blankets and Tritium Breeding

What Is a Blanket?

It is certainly not a thin, soft cover to keep the plasma warm. It is a thick, massive, complex structure that serves three major purposes: (1) capture the neutrons generated by fusion and convert their energy into heat, (2) produce the tritium to fuel the DT reaction, and (3) shield the superconducting magnets from the neutrons. The blanket is divided into modules for easier replacement. Figure 9.9 shows where the blanket is located inside a tokamak. In Fig. 9.9a, we see that the plasma first strikes the first wall (FW), which is also the front surface of the blanket. Then, the neutrons go into the blanket, where their energy is captured, and where the tritium breeding takes place. The heat is taken away by hot gas or liquid coolants to heat exchangers outside. Shielding material protects the vacuum walls and superconducting magnets from the heat and the neutrons. Figure 9.9b gives an idea of how the blanket surrounds the plasma and lies inside the vacuum. Outside the vacuum vessel are the magnetic coils. The Central Solenoid coil is critical, since there is not much room in the hole of the torus to fit this coil into. The symmetry axis of the torus is at the left. The entire machine fits inside a cryostat which insulates the magnet coils from the outside world, keeping them at superconducting temperatures.

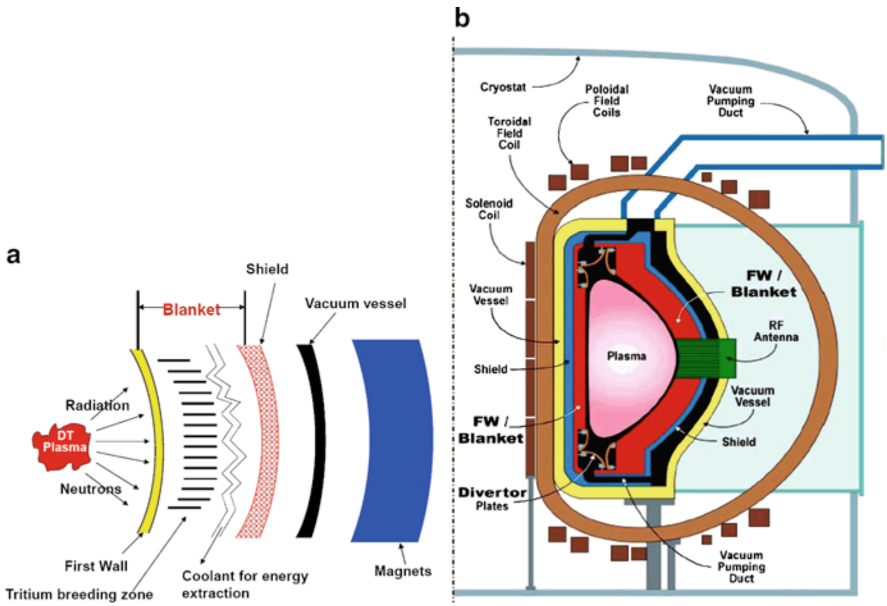

Fig. 9.9 (**a**) The order of the main layers in a tokamak, showing that the entire blanket must be inside the vacuum chamber. (**b**) General scheme of a tokamak's components, showing that the entire machine is inside a cryostat to keep the superconducting magnets cold [32]

In a reactor there could be hundreds of blanket modules, each weighing a ton. There are many ideas for blanket design, and ITER will have three ports available for test blanket modules (TBMs). There are six TBM proposals competing for these three spots [5].

The Role of Lithium

Deuterium and tritium are not the only fuels in fusion; lithium is needed for breeding tritium, which occurs only in minute amounts in nature. Lithium is an abundant element on earth, occurring in two isotopes, 92.6% Li^7 and 7.4% Li^6. (The superscript is the atomic weight, the total number of protons and neutrons in the nucleus.) Lithium-6 is the more useful one and can easily be enriched to 30–90% for use in a blanket. A 1,000-MW fusion plant will consume 50–150 kg of tritium a year, much more than can be supplied by other sources, such as fission reactors. To generate this amount of tritium in blankets, less than 300 kg of Li^6 will be needed by each reactor per year. About 10^{11} kg of lithium is available on land, and 10^{14} kg in the oceans. If all the world's energy is generated by fusion, the lithium will last 30 million years [6]. Deuterium will last even longer. There are 5×10^{16} kg of deuterium in the oceans, and at the rate of 100 kg per reactor per year, that will last 30 *billion* years! That's what we mean when we say that fusion is an infinite power source.

The way tritium is made from lithium-6 is shown in Fig. 9.10. The neutron, which started out at 14-MeV energy, has been slowed down by collisions with a moderator material and collides with a lithium nucleus, breaking it into an alpha (α) particle (helium nucleus) and a triton (tritium nucleus). Together, these carry the 4.8 MeVs of energy which is gained in splitting the lithium nucleus. This energy, as well as the neutron's energy, is transferred to a liquid or gas coolant and eventually transferred to steam for generating electricity. The n-Li^7 reaction is the same, except that a slow neutron is left over which can undergo another tritium-producing reaction. The n-Li^7 reaction works only with fast incoming neutrons, however.

The problem with this scheme is that not enough tritium is produced, since only 20–40% of the neutrons actually react with the lithium [3]. Some of the neutrons are lost through gaps in the blanket needed for plasma heating and measuring equipment. Some are lost by striking support structures instead of the lithium-bearing material, and a few are lost by passing through the whole blanket. To make up for this, there are fortunately neutron multipliers, mainly lead (Pb^{208}) and beryllium (Be^9). These can yield two neutrons for each incoming one. The reaction for beryllium is shown in Fig. 9.11.

Blankets will contain lithium, lead, beryllium, and a structural material; but the main problem is to cool them to take out all the heat that is the power output of the reactor. Blanket designs differ in the way they are cooled and in the form of lithium that is used. To show what is involved, we shall describe three of the leading proposals that have been worked out in detail.

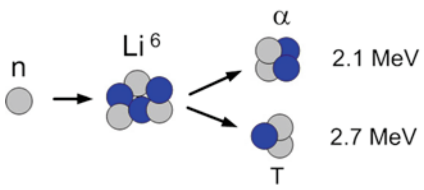

Fig. 9.10 The n-Li^6 breeding reaction, in which a neutron breaks a lithium-6 nucleus into an alpha particle (helium nucleus) and a triton (tritium nucleus). Protons are *blue* and neutrons are *gray*

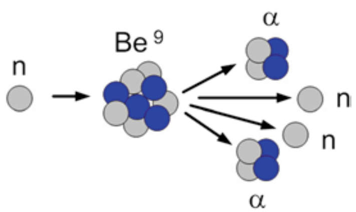

Fig. 9.11 Beryllium acts as a neutron multiplier, breaking up into two helium nuclei and two neutrons when joined by a neutron

Blanket Designs

The main coolants available are pressurized water, liquid metals, and helium. Water can be used only for near-term experiments. Reactors will probably need helium gas at a high temperature. The structural materials would be the same as those considered for the first wall: ferritic steels, vanadium alloys, or silicon carbide composites. The lithium can be in the form of solid pebbles of lithium ceramic, a liquid mixture of lead and lithium, or a molten salt called FLiBe [3]. Figure 9.12 shows how a TBM will be inserted into one of the ports in ITER.

The *helium-cooled ceramic breeder* (HCCB) uses solid material, with the beryllium multiplier and the lithium breeder in separate compartments. Figure 9.13 shows

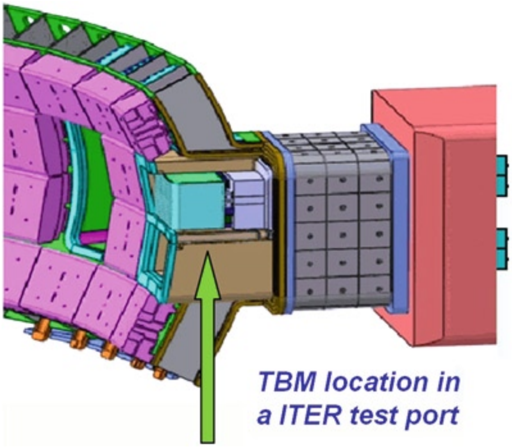

Fig. 9.12 Provision for insertion of test blanket modules in ITER, replacing part of the first wall [33]

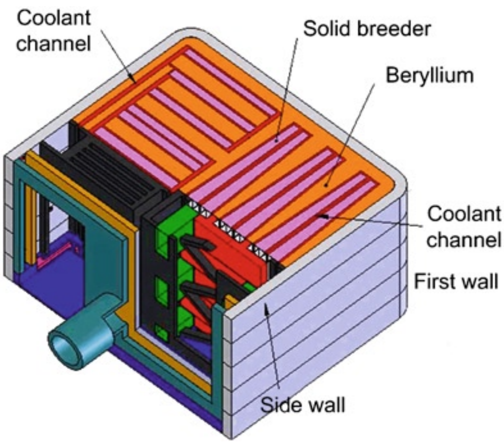

Fig. 9.13 Schematic of a helium-cooled ceramic breeder module [33]

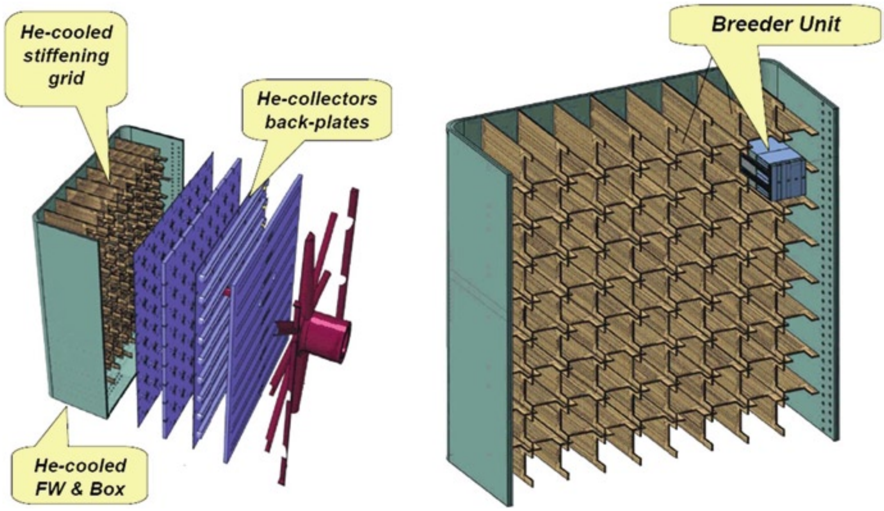

Fig. 9.14 Schematic of a large blanket module. The exploded view at the *left* shows several layers of supporting grids and coolant pipes which have been slid out of the box for clarity. The first wall (FW) is at the *left*. The view at the *right* shows the slots into which the submodules will be placed [3]

the parts of an HCCB module. The slabs containing the beryllium and the lithium ceramic are shown in red and blue. Between the slabs are cooling channels through which helium is pumped under 80 atmospheres of pressure [3]. The temperature of the helium can reach 500°C, and the breeder material can reach 900°C. Note that the front of the blanket is part of the first wall. In a reactor, a blanket module can be assembled from submodules, as shown in Fig. 9.14. The thickness of the blanket is about 50 cm and its width about 3 m.

The solid breeding material consists of ceramic pebbles of lithium orthosilicate (Li_4SiO_4), lithium metatitanate (Li_2TiO_3), or other similar materials. Techniques have been developed to manufacture identical spherical pebbles which can distribute themselves uniformly. The size should be small, less than 1 mm in diameter, to minimize the temperature difference across the radius so that the brittle spheres do not crack [7]. To extract the tritium, a flow of helium containing some deuterium (D_2) or hydrogen (H_2) is passed through the pebble bed, and the tritium (T_2) is carried out in the flow. The gases are then frozen and separated by distillation, since each has a different boiling point. The important thermal properties of a pebble bed have been measured [8].

A *helium-cooled lithium lead* (HCLL) blanket uses a molten alloy of lithium and lead called a eutectic. Meaning *easily melted* in Greek, a eutectic melts at a lower temperature than its constituents. The preferred eutectic is *Pb-17Li*, containing 17% lithium enriched to 90% Li^6. This melts at 234°C, compared with 328°C for lead and 181°C for lithium. In a blanket, the eutectic is heated from 400 to 660°C by the neutrons [3]. Since lead is a neutron multiplier like beryllium (Fig. 9.11), the multiplying and breeding are done in the same liquid. The submodules

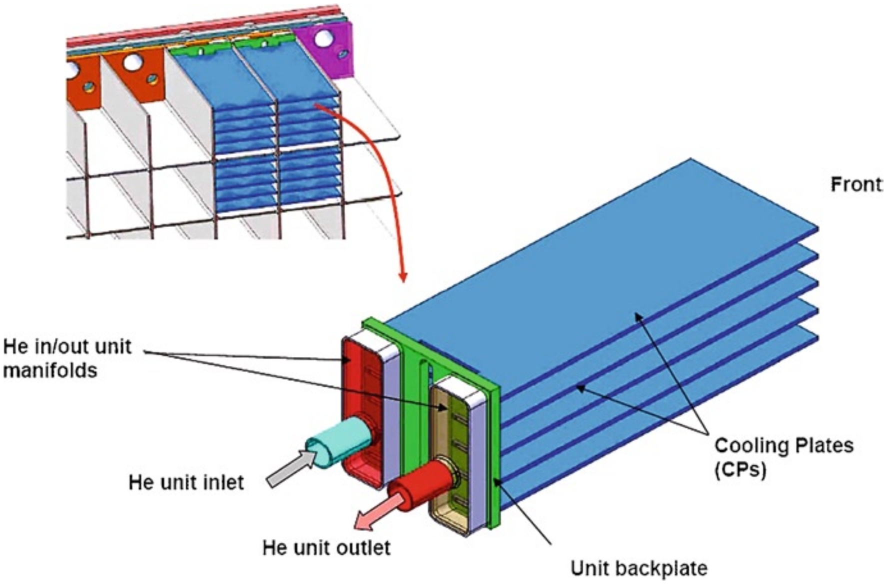

Fig. 9.15 Helium cooling arrangement in an HCLL blanket submodule [3]

in Fig. 9.14 will have circulating paths for the Pb-Li interspersed with channels for the helium coolant. The helium part is shown in Fig. 9.15, and the Pb-Li will go between the cooling plates. The tritium generated in the Pb-Li can be recovered by one of the two methods: permeation or bubbling. Hydrogen has a tendency to diffuse through walls, and tritium is just another form of hydrogen. Inside the blanket, tritium permeation into the helium coolant or other places where it does not belong is to be avoided. Outside the blanket, however, permeation windows can be made to allow hydrogen to go through and mix with a helium flow headed for a tritium separation facility. Alternatively, the Pb-Li can be formed into bubble columns where bubbles of helium capture the tritium in the liquid Pb-Li and carry it to the processing plant.

In earlier work, a molten salt called FLiBe, containing beryllium fluoride (BeF_2) and one or two parts of lithium fluoride (LiF) was proposed as a breeder fluid, but now Pb-Li is preferred. The work on FLiBe uncovered the problem of magnetohydrodynamic flow [9], which also applies to Pb-Li [10]. Both are electrically conducting fluids, and when these move inside a magnetic field, electric currents are generated in the fluid; and these currents react back on the magnetic field to produce a drag on the fluid motion. Considering how strong the magnetic fields are in a tokamak, this drag is a serious problem that increases the required pumping power. The drag is less if the flow goes along the magnetic field lines, but eventually the fluid has to cross the field lines to get out of the breeding region.

A *dual-cooled lithium lead* (DCLL) blanket uses both helium and the Pb-Li itself as coolants. This concept is shown in Fig. 9.16. Since Pb-Li is a liquid, it can be sent to its own heat exchanger and act as its own coolant. Helium is used to cool

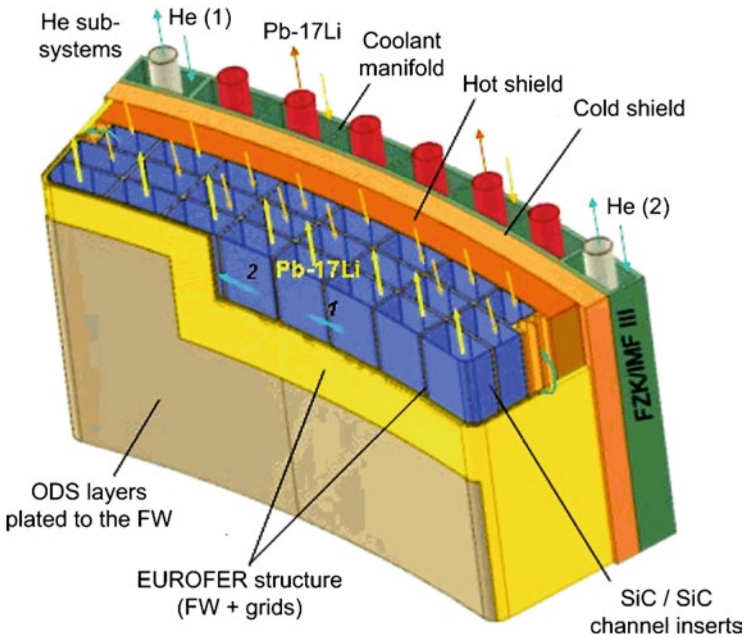

Fig. 9.16 Schematic of a dual-cooled lithium lead blanket module [34]. ODS, EUROFER, and SiC/SiC refer to high-temperature materials described under The First Wall and Other Materials

the first wall separately. The flow in the Pb-Li channels is shown in Fig. 9.17 for a case in which the magnetic field direction is into the paper. Computer models have been developed to describe the flow of the conducting liquid, including the buoyancy effect when the temperatures at the top and bottom are different. The eddies in the flow, as calculated, are shown in the inset. Since each module in a tokamak will be oriented at a different angle to the magnetic field, the structure of the flow, and hence the pressure drop, will be different at each location in the machine.

In advanced designs, the helium is eliminated, resulting in a self-cooled lithium lead breeding blanket, in which Pb-Li does all the cooling. It may take a lot of power to pump Pb-Li fast against the drag by the magnetic field. The possibility also depends on the development of the wonder-material SiC/SiC, which can operate at 1,000°C and contain a higher temperature fluid than other materials.

These blanket designs do not show all the auxiliary equipment necessary to operate the blanket. The roomful of pipes, heat exchangers, shields, and instruments for a single TBM in ITER is shown in Fig. 9.18. The blanket module itself is only the curved unit at the left, which forms part of the first wall.

Blankets for a full-scale reactor would have to satisfy many other requirements besides cooling and breeding. *Maintenance and operation* presents serious problems for a reactor designed to operate for over 25 years. The blanket material will have to be replaced many times during the life of the reactor. Solid breeders such as the

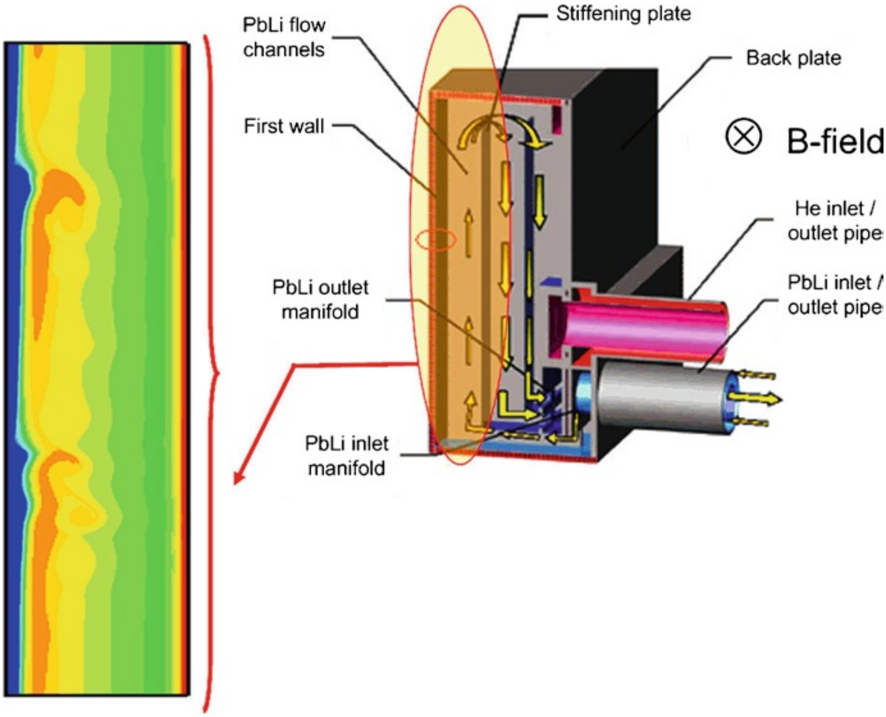

Fig. 9.17 Lead-lithium flow paths in a DCLL blanket submodule. The inset shows computer results for the eddy currents in one of the columns when the flow is perpendicular to the magnetic field [32]

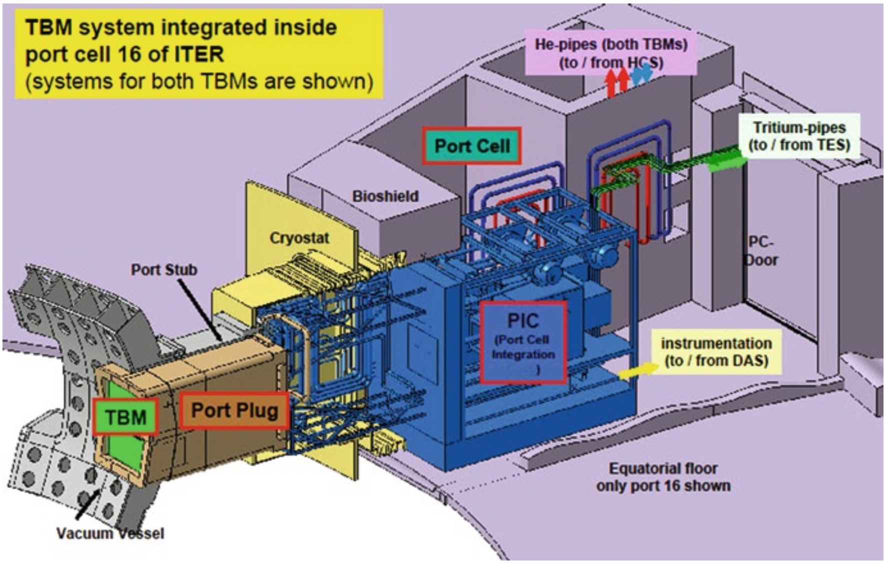

Fig. 9.18 Diagram of a proposed test blanket installation in ITER [6]

pebble-bed HCCB have to be physically removed to change the pebbles. In liquid blankets, the Pb-Li can be circulated outside the blanket and renewed without removing the blanket. Eventually, however, blankets will have to be replaced, requiring a shutdown. For easier replacement, banana-shaped blankets fitting the contour of the D-shaped plasma have been proposed. These would be lowered from the top of the tokamak during a shutdown, and all the connections to the blanket would have to come from the top. All this has to be done with remote handling, since there will be too much radioactivity for humans to work on the reactor.

Since the blankets are located inside the vacuum, they must be leak proof. Welds must be secure. Inside the blanket there are many interfaces between breeders and coolants, and a leak there would be impossible to fix without removing the blanket. There are also numerous joints in the pipes connecting the blanket to the world outside the vacuum tank. In 2008–2009, the Large Hadron Collider in Geneva suffered from a single leak in the liquid helium system which delayed the startup of the machine for over a year. In 2003, a single piece of loose foam brought down the shuttle Columbia, killing seven astronauts. Accidents happen, and extreme care must be taken in a tokamak reactor, where there are a million places where a leak can occur.

There are also safety issues in the case of an accident, including decay heat and radiotoxicity after shutdown [11]. Recycling and treatment of waste have also to be considered. However, these are not specific to blankets and will be covered in another section.

Tritium Management

Tritium Self-Sufficiency

The blanket designs shown above can barely breed enough tritium to keep a D–T reactor going. The tritium breeding ratio (TBR) is a measure of this. Each time a T fuses with a D in the plasma, one neutron is created. This neutron has to generate more than one T to re-inject into the plasma because there will be losses in the process. In addition, extra T's have to be stored to build up the inventory of tritium to run the reactor at a higher power or to fuel another fusion reactor. Only fusion can produce the enough tritium to build up its own industry.

The number of T's created in the blanket for each incoming neutron is the TBR. It has not been possible to design a blanket with a TBR larger than 1.15. That means that less than a 15% margin is available. The consequence is that tritium self-sufficiency can be achieved only after many years. The time is long because only a small percentage of the tritium injected into the plasma actually fuses with a D; most of it goes out the divertor and is recycled. This *fractional burnup* is only a few percent. Figure 9.19 shows calculations of how long it will take to double the tritium inventory. On the vertical scale, the TBR is plotted. The bottom portion, below TBR = 1.15, is what is possible. The horizontal axis shows the fractional burnup in percent. The curve labeled 1 year shows that it is not possible to double the tritium

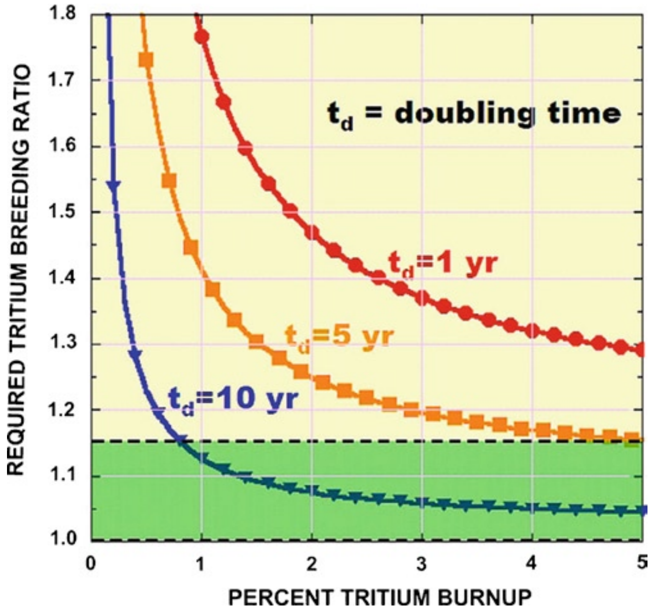

Fig. 9.19 Curves of the doubling time of the tritium inventory plotted against the TBR and fractional burnup of the tritium [32]

inventory in 1 year, since the curve never goes low enough to reach the feasible range of TBRs. The 5-year curve barely makes it if 5% burnup can be achieved. More likely, it will take almost ten years to double, and self-sufficiency can be achieved only after decades.[2]

In early tokamaks, before good divertors were developed, the fractional burnup was much larger, perhaps 30%, because of *recycling*. Ions of the plasma would hit the vacuum wall and recombine into neutral gas. This gas would go back into the plasma and be re-ionized and re-heated, thus being available again without having left the chamber. If modern divertors work well, however, ions are prevented from hitting the wall, thus preventing recycling. The ions are instead led to the divertor, where they recombine into gas and are pumped out before they can re-enter the plasma. In ITER, the fractional burnup is expected to be only 0.3%, which would be unacceptable for reactors [32]. Since burnup depends on the triple product $Tn\tau$ discussed in Chap. 8, this is another indication of the large step between ITER and a working reactor.

A fission reactor can produce only 2–3 kg of tritium a year, and tritium decays by 5.5% per year, so it is continually being lost. It will take 10 kg of tritium just to get DEMO started. ITER itself will use up most of the tritium available in the world [32]. There is therefore some urgency to develop breeding blankets with higher TBRs.

Tritium Basics

As doubly heavy hydrogen, tritium has two extra neutrons, which do not sit well with a single proton. So tritium decays by emitting an electron, a process known as beta-decay. This loss of a negative charge changes one of the neutrons into a positively charged proton and converts tritium into helium-3, a helium isotope with two protons and a single neutron instead of the usual two. This decay makes tritium radioactive, and it has to be handled carefully in a fusion plant.

Fortunately, the radioactivity is mild. The electron that is emitted has very low energy, about 19 keV. It cannot penetrate the skin, and even in air can go only 6 mm (1/4 in.) [12]. However, it can be harmful if ingested and must be carefully kept out of the water supply. Unlike fission products, tritium has a short half-life of only 12.3 years. This means that 5.47% of it decays into harmless helium each year. Because of its short life, very little tritium exists naturally. Cosmic rays make about 200 g of tritium a year, and there are only about 4 kg of natural tritium at any one time in the earth's atmosphere. Man-made tritium raises this to about 40 kg. Compared with this, it will take 1 kg of tritium just to get ITER running on DT, and a reactor may use up 100 kg per year.

The Tritium Fuel Cycle

One of the most complex technological tasks is to manage the supply of tritium. Tritium is injected into the plasma as fuel. It leaves the plasma through the vacuum pumps, most of it going through the divertor. It is generated in the breeding blankets and has to be captured and purified. It is also a contaminant in the liquids and other materials that leave the reactor and has to be removed from them. Excess tritium has to be stored safely for future use in raising the power of the reactor or starting up other reactors. Figure 9.20 shows a simplified diagram of these paths.

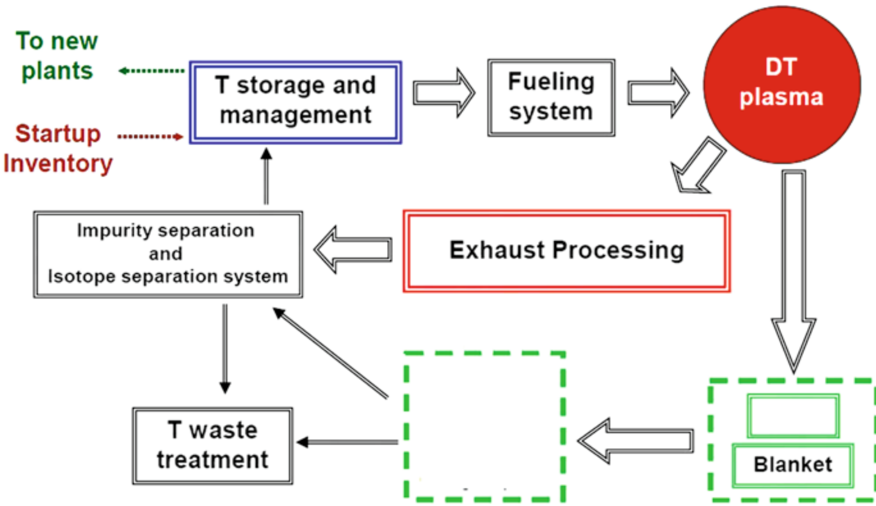

Fig. 9.20 Simplified diagram of tritium fuel cycle [32]

Tritium leaves the tokamak in two paths – either through the vacuum pumps, including those pumping the divertor, or through the first wall (FW) and the blanket. The vacuum exhaust goes directly to an isotope separation system which saves the T_2, D_2, and He and removes the impurities. Pure T_2 is sent directly to Tritium Storage and Management. The tritium generated in the blanket goes first to a tritium processing plant to remove it from the breeder materials, and then to isotope separation. Material contaminated with irremovable tritium from both streams then goes the Tritium Waste Management. The fueling system receives recovered tritium from the two paths as well as from storage or from external sources. The fueling system then injects tritium and deuterium into the plasma. Deuterium is cheap and safe and does not have to be parsimoniously recovered.

The vacuum in the torus is maintained by cryo-pumps [13]. These are porous carbon surfaces cooled by liquid helium to 5 K; that is, 5° above absolute zero, the latter being −273°C or −459°F. At that temperature, all gases except helium are condensed and stuck to the cryogenic surfaces. To release hydrogen, including tritium, the cryo-pumps are periodically heated to about 90°K, and this gas is sent to the isotope separation system. To release all the captured gases, the pumps are raised to room temperature.

Fueling is done by injecting frozen pellets of tritium and deuterium at sufficient velocity to reach the center of the plasma. This is much more efficient than injecting DT gas at the boundary, since the gas will be ionized at the surface and will not reach the interior. There is some loss of tritium in the process, and this will appear in the pumping system. The plasma is heated mainly by neutral beam injection (NBI), the beams consisting of deuterium and tritium. This system will have its own system of tritium management.

Isotope separation is done by freezing the gases to liquid helium temperatures and selective warming in four interlinked distillation columns [13]. The tritium processing plant in ITER is a large seven-story building [12]. In addition, all water in the ITER installation and all air from buildings have to pass through a detritiation plant to remove the tritium. Water released back into the environment is pure H_2O, and hydrogen released into the air is pure protium (H_2). Tritium has to be stored until it is used. This is done in metal-hydride getter beds, each capable of holding 100 g of tritium [13]. Zirconium–cobalt (ZrCo) absorbs T_2 to form $ZrCoT_3$. The reaction is reversible upon heating to release the T_2. Although techniques for tritium containment are well established in the fission industry, the amount of tritium in fusion is orders of magnitude larger. There has been no experience so far on such a large scale.

Superconducting Magnets

Introduction

The dominant features of a tokamak or any other magnetic bottle are the heavy coils that generate the large magnetic field used to confine the plasma. Until recently, all tokamaks had magnet coils made of copper, which conducts electricity better than any other metal except silver. Even so, it takes a lot of energy to drive megamperes of current through copper coils, and fusion reactors will have to use superconducting coils. Superconductors have zero resistivity, and once the current has been started in them, it will keep going almost forever. The hitch is that superconductors have to be cooled below 4.2 K with liquid helium. A cryogenic plant has to be built to supply the liquid helium, and the magnet coils (and hence the whole machine) have to be enclosed in a cryostat to insulate them from room temperature. The good news is that this technology is well developed and is *not* one of the serious obstacles to fusion power. In 1986, the world's largest superconducting magnet, the MFTF (mirror fusion test facility), was completed at the Lawrence Livermore Laboratory in California. It was a different type of magnetic bottle that we will describe in Chap. 10. However, the program was almost immediately canceled by the Reagan administration in favor of the tokamak because the USA could not afford to follow two expensive paths to fusion. The MFTF was so large that for a while it became a museum that one could walk through. Currently, three superconducting tokamaks are in operation: the Tore Supra in France, the EAST (Experimental Advanced Superconducting Tokamak) in Hefei, China, and K-STAR, in Daejon, Korea. Soon to join them is an upgrade to Japan's JT-60U (Fig. 8.6) called JT-60SA. In addition, the Large Helical Device, a superconducting stellarator-type machine, has been operating for two decades in Japan. ITER will, of course, have superconducting magnets.

Two superconducting materials are available on a large scale: niobium–titanium (NbTi) and niobium–tin (Nb_3Sn). NbTi is cheaper and easier to make, but it loses its superconductivity above 8 T. A tesla is a large unit of magnetic field equal to 10,000 G, the old unit. Common magnets rarely go above 0.1 T, but some magnetic resonance imaging (MRI) machines in medicine can go up to 1.5 T. The earth's magnetic field is only about 0.5 G or 0.00005 T. In ITER, fields up to 13.5 T are needed, so some coils are made of Nb_3Sn, and others (for lower fields) are made of NbTi. The dividing line is around 5 T [14]. Superconducting cables are complicated to make because they have to be made of a thousand thin strands. This is because the current in superconductors flows only on the surface, and thin strands have large surface areas compared to their volumes. Also, the cables have to be bendable.

ITER's Magnet Coils

Figure 9.21 shows what a niobium-tin cable looks like inside. There are over 1,000 strands in six bundles. At the center is a helix making room for the pipe that carries the liquid helium. The outer casing is a stainless steel jacket 37.5 mm (1.5 in.) in diameter. This cable, designed for the toroidal field coils of ITER, can carry 80 kA at 9.7 T. Each strand is about 0.8 mm in diameter and consists of a Nb_3Sn filament sheathed with chromium and covered with about as much copper as Nb_3Sn. The copper is necessary to mitigate quenches. A quench occurs when part of the superconductor goes normal, losing its superconductivity because of over-heating or over-current. Huge voltages would build up as the current tries to force

Fig. 9.21 Construction of a niobium-tin cable. One of the bundles has been exploded to show the strands [14]

its way through a normal conductor with resistance, and there could be an explosion. Copper can make this a gentler accident. The complexity of superconducting cables is bad enough, but to wind them into magnetic coils means that each cable has to be over 1.5 km (a mile) long.

A tokamak has many different kinds of magnet coils, and each requires a different design. Some of these can be seen in Fig. 9.22. The toroidal field (TF) coils are the large D-shaped coils. They operate up to 6 T and are the heaviest ones. Transporting them to the ITER site requires special barges, trucks, and roads. The large, horizontal ones encircling the machine are the poloidal field (PF) coils, which give the field lines their twist and shape the plasma. Because of their size, they cannot be transported and must be wound on site. The coil winding building at the ITER will be 253 m long, 46 m wide, and 19 m high.[3] A critical component is the central solenoid (CS), seen inside the hole in the torus. There is very little space there, and most of it is taken up by the interior blanket modules. This coil is the other half of the PF system that shapes the plasma and drives the tokamak current. The CS is 13 m tall and 4.3 m in diameter, weighing 1,000 tons. It also produces the highest field of 13.5 T. Figure 9.23 shows a test section of it that has been made.

There are smaller coils besides these main coils, but the difficult part is to join the superconductors to their feeds. Current is fed into the coils from normal-conducting cables, and then a superconducting switch is turned on so that the

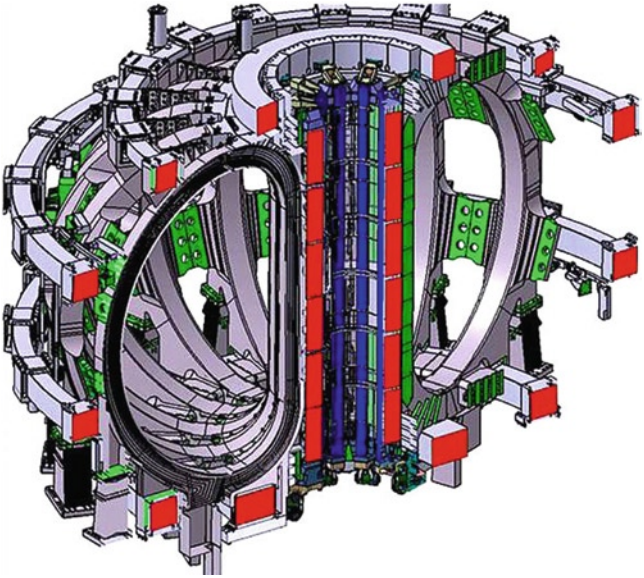

Fig. 9.22 Drawing of the magnetic coils in ITER (ITER Newsline Nos. 114 and 122 (2010). http://www.iter.org/newsline/)

Fig. 9.23 A test section of the Central Solenoid for ITER [14]

current flows only in the superconductors and the feed cables can be disconnected. These junctions are very complicated, especially since the current has to go through the wall of the cryostat from room temperature to 4 K. Almost all the nations supporting ITER participate in designing and producing the magnet system. Some make the NbTi and Nb_3Sn materials. Some make it into strands and cables. Some wind the cables into coils. And some make the feed cables and the junctions. The technology has already been developed for smaller tokamaks, and the steps to ITER, DEMO, and reactor are only matters of scale.

The Supply of Helium [4]

Helium is not a rare gas if we can afford to fill the world's balloons with it. Actually, balloons account for only 16% of helium use. Cooling of semiconductors accounts for 33%, and the rest is used for industrial and scientific purposes. The atmosphere contains four billion tonnes of He, but it is not economical to extract it by cryo-distillation. Most of our helium comes from natural gas as a by-product. Thus, helium comes from fossil fuels and will be depleted in several decades along with natural gas, as discussed in Chap. 2. In this chapter, we have seen how critically fusion reactors, as envisioned today, will depend on helium in both extremely hot and extremely cold places. In the first wall and blankets, gaseous helium is used as a high-temperature coolant. The vacuum system uses liquid helium to cool the cryo-pumps. In the magnet system, liquid helium is what produces superconductivity. It is a closed system, but there are leaks. It is estimated that ITER will lose 48 tonnes of helium a year, about 0.15% of the world's current consumption. But if eventually fusion produces a third of the world's power, those reactors would need the world's supply of helium for a whole year just to start up [4]. At some point the helium losses, say, 10% of the inventory, would exceed what comes from natural gas. You will remember that helium is one of the products of the D–T reaction. At only a few percent burnup, however, this "ash" is a negligible contribution to the total demand. Helium is needed in other industries as well; for instance, in medical equipment. The shortage is so acute that a rationing system was proposed in the USA in 2010.

High-Temperature Superconductors

In 1986, compounds were discovered that became superconducting at a critical temperature as high as 30 K. Since then, research to find better materials has been intense. The goal was to get the critical temperature above 77 K, the point at which nitrogen becomes liquid. Liquid nitrogen is much, much cheaper and easier to produce than liquid helium, which is liquid below 4 K. The 73°C difference between 77 and 4 K does not seem much. We encounter such a change every time we boil a cup of coffee. However, since one can never go below absolute zero, it is the distance from absolute zero that is important. Seventy-seven kelvin is 19 times farther from 0 K than is 4 K; and, of course, there is no shortage of nitrogen. The goal has already been achieved; three superconductors have been found that work at liquid nitrogen temperatures. The record as of 2009 is 135 K, well above 77 K. Typically, the compound is complicated: $HgBa_2Ca_2Cu_3O_x$. Until searches can be made by computer, finding new compounds will be slow; but it is a reasonable expectation that large-scale production of a high-temperature superconductor will be possible by the time DEMO is built. Maybe a room-temperature superconductor will have been found by that time. The machine would be much simpler and cheaper.

Plasma Heating and Current Drive

Introduction

Bringing the plasma up to fusion temperatures is done with the injection of neutral atoms and the excitation of different types of plasma waves. In addition, waves are also used to drive the plasma current without using transformers – so-called noninductive current drive. There are many physics problems involved in these processes. Neutral beams also fuel the plasma and give it rotational velocity. Waves not only heat the plasma and drive its current but are also used to change local conditions inside the plasma and shape the current profile. In this chapter, we are concerned with technology and therefore concentrate on the hardware and discuss only the main types of waves that can be used.

Neutral Beam Injection (NBI)

One of the aims of ITER is to reach ignition, when the alpha particles generated by the D–T reaction are able to keep the plasma hot. To get to this point, however, immense power has to be injected to raise the temperature to the order of 50 keV (500,000,000°). This is done mainly with NBI. ITER will have 33 MW of NBI. The injectors, three or four of them, are usually the largest appendages sticking out of the tokamak. In the first stage, deuterium atoms are given an extra electron to produce negative ions. Once charged, the ions can be accelerated electrostatically. Before entering the tokamak, the negative ions go through a little gas, which strips off the extra electron, restoring the atom to a neutral state. Being neutral, the atom is not affected by the magnetic field and can go well into the plasma until it is ionized by the electrons in the plasma. How far it goes depends on its energy. All large tokamaks use NBI, which is a well-established technology; but since ITER is so large, neutral beams of 1 MeV energy are needed to get to the center. NBI technology for 1 MeV has not yet been developed [15].

Ion Cyclotron Resonance Heating (ICRH)

This method heats ions by pushing them with a rotating electric field whose direction follows the ions' cyclotron motion as they revolve in their nearly circular Larmor orbits. It is sometimes more efficient to heat a minority species, such as helium-3 rather than deuterium or tritium, because of the way the energy is coupled into the plasma. The cyclotron frequency depends on the magnetic field strength, so the applied electric field has to be of a specific frequency, depending on magnetic field at the location where the ions are to be heated. In ITER, this frequency is in the range around 50 MHz. This is too low a frequency to be transmitted through a pipe, so an antenna has to be placed inside the vacuum chamber. The antenna is outside the field lines leading to the divertor (see Fig. 9.4), but it is so close to the plasma that it will be bombarded by ions. These ions will sputter antenna material into the plasma, and such contamination usually cools the plasma. ITER is to have 20 MW of ion cyclotron heating. The power is not the main problem; the problem here is to design antennas which will not affect the plasma deleteriously.

Electron Cyclotron Resonance Heating (ECRH)

In principle, what is done to the ions can also be done to the electrons, but the technology is entirely different. The electrons' cyclotron frequency is in the gigahertz range, and huge microwave generators are required. The power or current input can be deposited accurately at specific places inside the torus by adjusting the microwave frequency to match the magnetic field at those places. Since microwaves are carried through waveguides, which are specially sized and shaped pipes, they can be injected through holes in the first wall and do not require an antenna inside the vacuum chamber. The bad news is that electron cyclotron waves cannot penetrate into the plasma from the outside of the torus. A property of these waves is that they must be injected from a high magnetic field into a lower magnetic field. Since the magnetic field is highest in the hole of the torus, the launching waveguide must be located in the cramped space also occupied by the central solenoid and the inside blankets. Waves at twice the cyclotron frequency, which also resonate with the electrons' gyrations, can get in from the outer, weak-field side; but the higher frequency is more difficult to generate.

The electron cyclotron heating system in ITER calls for 20 MW of power at 170 GHz. This frequency corresponds to the cyclotron frequency at 6.0 T (60,000 G), high enough to cover ITER's magnetic field of 5.5 T at the inside radius. Although we use microwaves in everyday life, 20 MW at 170 GHz is an entirely different matter. A microwave oven puts out 1 kW at 2.45 GHz using a magnetron so small that we are not aware of its presence. Powerful microwaves are generated by gyrotrons, which work by running ECRH in reverse. In a gyrotron, an energetic electron beam is first produced. It is then injected into a magnetic field, so that the electrons undergo cyclotron gyrations. In doing so, they emit microwaves at harmonics of the cyclotron frequency which are then channeled into a waveguide leading to the tokamak. The microwaves get their energies from the electron beam, which loses part of its energy. In experiment, the remaining energy is captured in a beam dump as heat. In advanced gyrotrons, the beam can, in principle, be re-injected so that its remaining energy can be re-used. Note that the electron beam in a gyrotron cannot be injected directly into a tokamak to heat it because the electrons cannot get through the magnetic field. In a gyrotron, the electrons are injected into the magnetic field from the *ends* of the field lines. A tokamak, of course, *has* no such ends; hence the need to convert kinetic energy into microwave radiation and then injecting the radiation instead of kinetic energy directly.

High-power gyrotron research began in St. Petersburg, Russia, decades ago. Those that can operate continuously for ITER are being developed in Japan, Germany, and the USA. So far, 1 MW at 170 GHz in a long pulse has been shown to be possible. Figure 9.24 shows the size of such a gyrotron. ITER will need 24 of these to produce the required ECRH power. Figure 9.25 shows a design of a 2-MW gyrotron with superconducting magnets.

Since the gyrotron has to be under vacuum and the waveguide is at atmospheric pressure, windows have to be used to isolate the waveguide from the vacuums at both ends. At present, the only material that can transmit the wave power at that frequency is synthetic diamond. Windows 10 cm (4 in.) in diameter have been made and tested for proper cooling. In a reactor, gyrotrons and their windows have to run continuously without failure for months or years between maintenance shutdowns. This constitutes a large step in engineering that has yet to be done.

Fig. 9.24 The gyrotron room at JAERI [35]. A 1-MW gyrotron is shown at the *left*. It is 3 m (10 feet) high and covered with magnetic coils

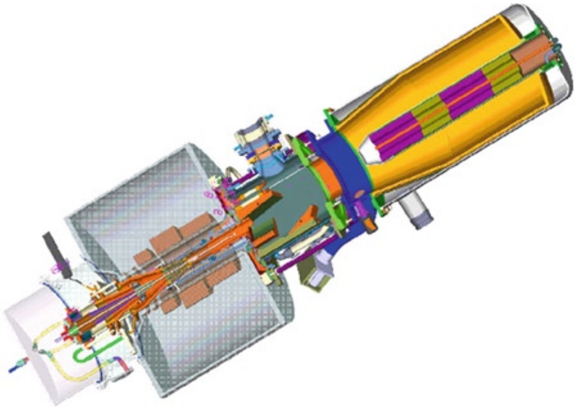

Fig. 9.25 Design of a 2-MW, 170-GHz superconducting gyrotron being developed in Germany [6]

Lower-Hybrid Heating (LHH)

A third type of wave that can be used for heating and current drive is the so-called lower-hybrid wave. This wave is particularly useful for current drive because it can control the current profile near the outside of the plasma. The lower-hybrid frequency lies between the cyclotron frequencies of the ions and electrons, or about 5 GHz in ITER. Klystrons are used to generate frequencies in this range. The wave has a long wavelength in the direction of the magnetic field, so to launch it requires a large "grill," meters in size, as shown in Fig. 9.26. Each of the openings is a waveguide fed by one or several klystrons, each with its own vacuum window. The phase of the wave emanating from each wave-guide is set so that the total grill, including some dummy waveguides, forms the wave that deposits its energy in the right place. Since the launcher lies close to the plasma surface, its materials must sustain the heat and neutron damage that that implies.

In summary, the physics of auxiliary heating and current drive is well understood, but the engineering of the power supplies and the wave launchers present some difficult problems.

Remaining Physics Problems

The ITER machine is an experiment large enough to require an international consortium. Its mission is to achieve a burning plasma, one in which the alpha particles produced by the D–T reaction can maintain the plasma's temperature without external heating. At this stage of construction, not all physics problems have been solved, though they may be solved by the time construction is finished. We hope that these problems will be solved in time for DEMO. However, the physics does not have to be completely understood for something to work. Books have been written on the physics of tennis, baseball, sailing, and even pizza. Sometimes, it is easier just to get on with it.

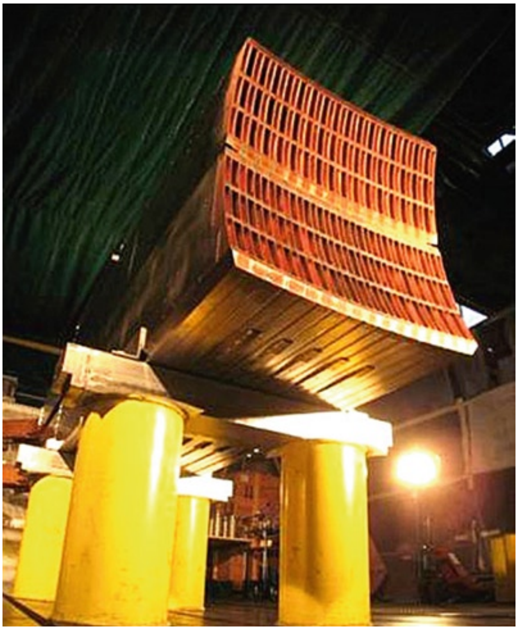

Fig. 9.26 A lower-hybrid wave launcher of the type designed for ITER but one-fourth the size [36]

Edge-Localized Modes

Edge-localized modes (ELMs) were described in Chap. 8. They are instabilities of the H-mode pedestal which can release plasma suddenly to the wall. Although most of these particles should flow to the divertor, the sudden burst of heat can erode and damage the divertor's surfaces. The H-mode pedestal constrains one-third of the plasma's energy, and 20% of this or as much as 20 MJ can be dumped into the divertor in a fraction of a second [16]. The preferred method to suppress ELMs is to impose a rippled magnetic field at the surface of the plasma, near the pedestal. The idea is to break up instabilities that tend to be aligned with the magnetic field. The pattern of currents in the ELM coils can be varied slowly to follow changes in the magnetic field lines. This method has been tested in the DIII-D tokamak in San Diego, California, and thorough calculations have been made to design the sizes and spacings of the coils for ITER [17]. A panel of ELM coils is shown in Fig. 9.27. Figure 9.28 shows what the surface of ITER will look like with these coils installed. It will take 2.6 MW of power to drive these coils. Being in-vessel components, the coils have to withstand intense heat and neutron bombardment. In ITER, the coils are protected from the plasma by a 50-cm thick, water-cooled, nonbreeding blanket whose only function is to attenuate the neutrons.[4]

In DEMO, there would be no place for ELM coils, since breeding blankets have to cover the machine to capture as many neutrons as possible. Locating the coils behind the blanket would probably be too far. ELM coils are *ad hoc*, temporary

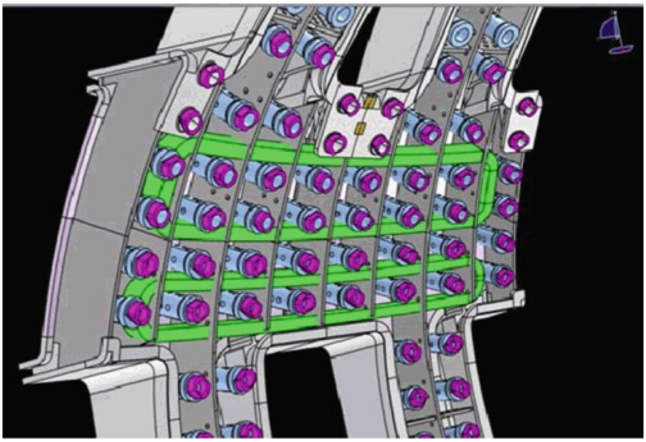

Fig. 9.27 A panel of ELM-suppression coils for ITER [6]

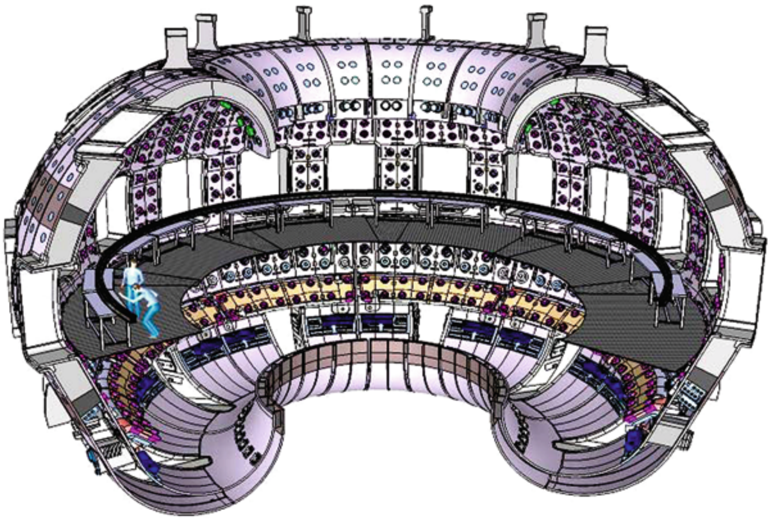

Fig. 9.28 Drawing of ELM coils installed in ITER [29]. The scale is shown by the human figures at the *left*

solutions not included in the original design of ITER since the problem had not yet arisen. The physics of ELMs has to be understood better to find passive methods for their control, but there is time to do this.

Once the ELM coils have been installed, they can also be used for other purposes. By applying a small current at a low frequency like 50 Hz, a weak instability called the RWM can be controlled. A differently spaced DC current can also be added to help prevent disruptions (described in detail in the next section).

Disruptions

As shown in Chap. 8, disruptions are disasters. Magnetic containment is suddenly lost, and the plasma drifts vertically into the walls, depositing all its thermal energy. The tokamak current tries to keep itself going as the plasma goes away, so very high voltages are generated. Runaway electrons of MeV energies are created by the high voltages, and these electrons crash into the walls, generating high-energy X-rays. The plasma current is used to generate the poloidal magnetic field, and as this field decays with the current, large forces are applied to the magnetic coils and conducting parts of the tokamak structure. The entire energy in the plasma, magnetic field, and tokamak current is something like 500 MJ, and in a disruption this is all dumped into the structure of ITER in 1/30th of a second [18]. This is like an explosion of 120 kg (260 lbs.) of TNT. Disruptions are expected in ITER, and its parts are designed to withstand them. Disruptions have to be eliminated in reactors, which would be so heavily damaged as to require lengthy shutdowns for repair.

There is a possible scenario of how a change in the magnetic structure of the tokamak discharge, such as a coalescing of magnetic islands, can cause a disruption. It has been confirmed in experiment that staying well below the known stability limits, such as the density limit, can avoid disruptions. A reactor, however, needs to operate at the highest level to lower the cost of electricity (COE). Since a disruption is now known to be a vertical displacement of the plasma, there are ideas on stopping these displacements with a coil or coils inside the chamber. Such a coil is included in Fig. 9.27. Though it is not possible to stop a disruption once it starts, there are ways to mitigate the damage. Disruptions have magnetic precursors which can be detected, and fast action can be taken. Injection of liquid jets or solid pellets of a frozen gas have been tried, but these have led to creation of too many runaway electrons. A large puff of a gas like argon can be driven well into the plasma, be ionized into high-Z ions, and increase the resistivity so that the current dies more gently. Fast gas valves have been developed for this purpose. There is then a smaller tendency to induce currents elsewhere, lower forces on the structure, and fewer runaway electrons. After a disruption, there is only gas left in the vacuum chamber. This has to be pumped out and the discharge started all over again.

Alfvén Wave Instabilities

In a burning plasma, 3.5-MeV alpha particles are generated, and as they cool down they transfer their energy to the plasma, keeping it hot. Before they become thermalized, however, the alphas are in the form of beams streaming along the magnetic field lines, and beams can excite instabilities. To do this, the velocity of the beam has to coincide with the velocity of a wave in the plasma; and the synchronism causes the beam energy to be transferred to the wave. The wave can become so strong that it disrupts the plasma. There is a plasma wave called the Alfvén wave that travels along the B-field and can have just the right velocity to match that of the alpha-particle beam. The danger that this can happen can be predicted precisely by theory [19], but whether it will actually happen or not depends on the details. ITER will be the first machine that can test for Alfvén wave instabilities in a D–T plasma. If these turn out to be important, their avoidance is a physics problem that needs to be solved.

Operating a Fusion Reactor

Startup, Ramp-Down, and Steady-State Operation

Turning on the power in a large tokamak is not an easy task. The vacuum system, the cryogenic system, discharge-cleaning of the walls, the magnetic field system, the tokamak current drive, and the various plasma heating systems, and various auxiliary systems have to be started up in sequence. Operators have learned by experience how to do this in large tokamaks. The plasma has to be maintained stably while it is being heated and while the current is being increased in synchronism with the toroidal magnetic field. Each power supply has to be ramped up at a certain time at a certain rate. Turning the discharge off also requires careful ramp-down of each system. Only after a good routine has been found can automatic controls take over.

All present tokamaks run in pulses, not continuously. Even if the pulses last for minutes or an hour, they will not uncover problems that will arise with truly steady-state operation. In the 1980s, a machine called the ELMO Bumpy Torus was run at the Oak Ridge National Laboratory. Though the magnetic configuration never caught on, the machine was run in steady-state and revealed problems that are not seen in pulsed machines. The Tore Supra tokamak in Cadarache, France, near the ITER site, has been gathering information on long-pulse operation for 20 years [20]. It is a large tokamak with high magnetic field, large current, and powerful heating. The first wall is water-cooled boronized carbon. In a deuterium plasma, the retention of deuterium by the carbon was found to be significant. This is one reason for rejecting carbon as a wall material. Damage to the ICRH antennas was noted. Electrical faults in the magnet system were found to limit the length of discharges. It was found that turning the lower-hybrid power on slowly greatly alleviated this problem. Water leaks were found to occur 1.7 times per year. The frequency of disruptions was also recorded. These were found to be caused mainly by the flaking of carbon off the walls after many days of operation. Pulses lasting 1 or 2 seconds were possible with transformer-driven currents, but with the addition of lower-hybrid current drive, 6-min pulses with 3 MW of lower-hybrid heating (LHH) were achieved in 2007. At the 2-MW level, 150 consecutive 2-min discharges could be routinely produced [21]. These are the types of problems that will be encountered when ITER is operated in continuous mode.

Maintaining the Current Profile

Advanced tokamaks utilize reversed shear and internal transport barriers for enhanced plasma confinement. These require precise shaping of the safety factor q (see Chap. 8), which determines how the twist of the magnetic field lines changes across the radius. The shape of the $q(r)$ curve controls the stability and loss rate of the plasma. Since the twist is determined by the poloidal field created by the plasma current, this current has to be shaped in a particular way. Some of the current is naturally produced by the bootstrap effect (Chap. 9); the rest has to be driven by lower-hybrid and electron cyclotron current drive. The blue curve in Fig. 9.29 shows an example of a $q(r)$ curve which stays above $q = 2$ and gives reverse shear. The red curve shows the auxiliary current needed to produce this $q(r)$. Only precise control of the localized heating can produce this current profile. As the plasma starts up, the currents will be changing, and the power supplies will have to be programmed to keep the current in a stable shape.

Remote Handling

Anytime tritium or deuterium is introduced into a magnetic bottle, the wall materials will become radioactive due to neutron bombardment. It will be impossible for humans to go inside the machine or even come close to it. Robots will be used to replace parts such as blanket modules, to fix leaks and make other repairs, and to examine the interior of the chamber during shutdowns. The robotic equipment itself will be exposed to neutrons. Such remote handling has been used successfully in

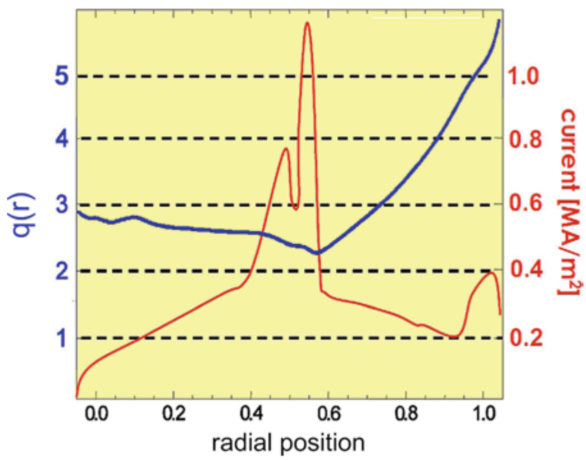

Fig. 9.29 Example of the variation of the safety factor $q(r)$ across the minor diameter of an advanced tokamak plasma (*blue*), and the plasma current distribution required to produce it (*red*) [37]

the TFTR machine at Princeton (Fig. 8.3) and the JET in England (Fig. 8.4), both of which have used DT fuel. Robots can weld joints by remote control. The first experiments in ITER will use hydrogen or helium, which produce no radioactivity. Later, deuterium experiments will give a small amount of radioactivity. In the next stage, tritium will be used; and the machine will become very "hot." ITER is much larger than TFTR or JET, and the components to be moved will be large and heavy. Remote handling is expensive and inconvenient, but it does not seem to be a technological barrier.

Fusion Development Facilities

The engineering of a fusion reactor will require solution of a number of serious technological problems, as we have seen above. ITER will take decades to build and operate, and it is not designed to solve many of these problems. It is therefore prudent to build smaller machines specially designed for technology development so that this work can proceed in parallel with ITER. Many proposals have been made for a fusion development facility (FDF). A few of these will be described here.

IFMIF: International Fusion Materials Irradiation Facility

A favorite proposal of the European Union, together with Japan, is the IFMIF, a large linear accelerator that has been in the planning stage for 16 years. A diagram of it is shown in Fig. 9.30. As you can see, this is a large installation. The accelerator occupies a building of several hundred meters in length. It is designed to produce neutrons with energies matching those that would enter a tokamak blanket. This is done by accelerating to 40 MeV a beam of deuterons onto a target of liquid lithium. Reactions like the reverse of that in Fig. 9.10 would occur: a deuteron on lithium-6 would produce beryllium and a neutron, and a deuteron on lithium-7 would produce beryllium and two neutrons. The neutrons would then be used to bombard different materials to see how they stand up.

The key parameters for assessing radiation damage are neutron flux, neutron fluence, and dpa. Flux is how many neutrons per second go through each square meter. Fluence is how many have gone through the area during the whole life of the material. Dpa measures the damage, either per year or for the whole life. The flux produced by IFMIF is comparable to that expected in ITER, and about four times less than that in DEMO. The dpa per year in IFMIF is comparable to that in DEMO (about 30) and much larger than at in ITER.[5] The fluence cannot compare with that in DEMO, but could duplicate that in the limited life of ITER.

The IFMIF will cost about $700M [22]. It has been severely criticized because only small samples, a few square centimeters in size, can be tested. This is entirely inadequate to test the large components of ITER and DEMO, especially the blanket modules.

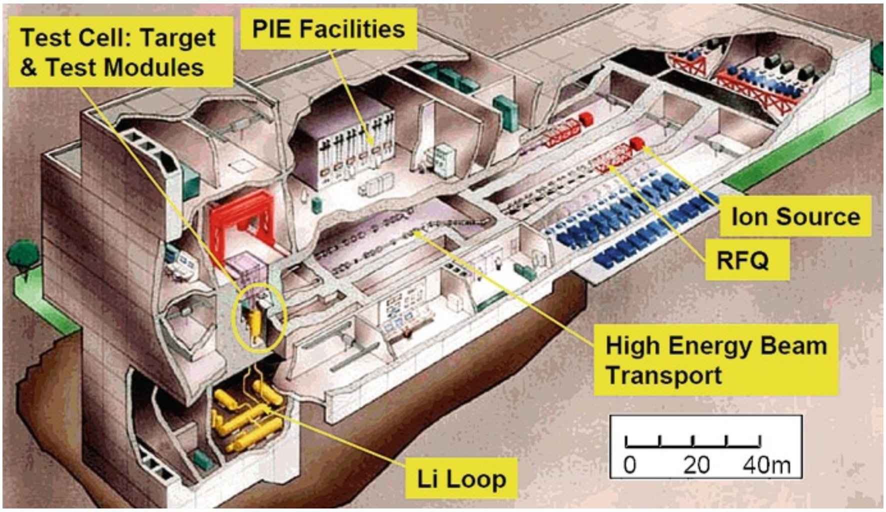

Fig. 9.30 Diagram of the International Fusion Materials Irradiation Facility [A. Möslang (Karlsruhe), *Strategy of Fusion Materials Development and the Intense Neutron Source IFMIF*]

Fusion Ignition Tokamaks

Proposals to build small but powerful tokamaks to test burning plasmas were made well before ITER. In the late 1980s, a Compact Ignition Tokamak was initiated in the USA, but was soon canceled. In 1999, Dale Meade at Princeton designed a 10-T, 2-m diameter tokamak call Fusion Ignition Research Experiment (FIRE), but this was never funded. These early ideas were based on the hope that very high magnetic fields produced without superconductivity could be used to achieve ignition on a small scale. This philosophy, promulgated by Bruno Coppi at the Massachusetts Institute of Technology, resulted in the Alcator tokamaks at M.I.T. and the Ignitor in Italy. In 2010, Italy and Russia signed an agreement to build a 13-T Ignitor-type tokamak to study burning plasma physics before ITER is finished. These small, pulsed machines cannot expose the steady-state problems that ITER will face. Engineering problems such as tritium breeding and plasma exhaust can be studied only with sufficient neutron flux. There are several proposals for large machines designed specifically for problems not tackled by ITER which will run simultaneously with ITER. None of these has been funded so far.

High-Volume Neutron Source

In 1995, noting the inadequacy of the IFMIF for blanket development, an international team headed by Abdou [23] proposed a high-volume plasma-based neutron source. A tokamak, naturally, was the best choice for a neutron source that

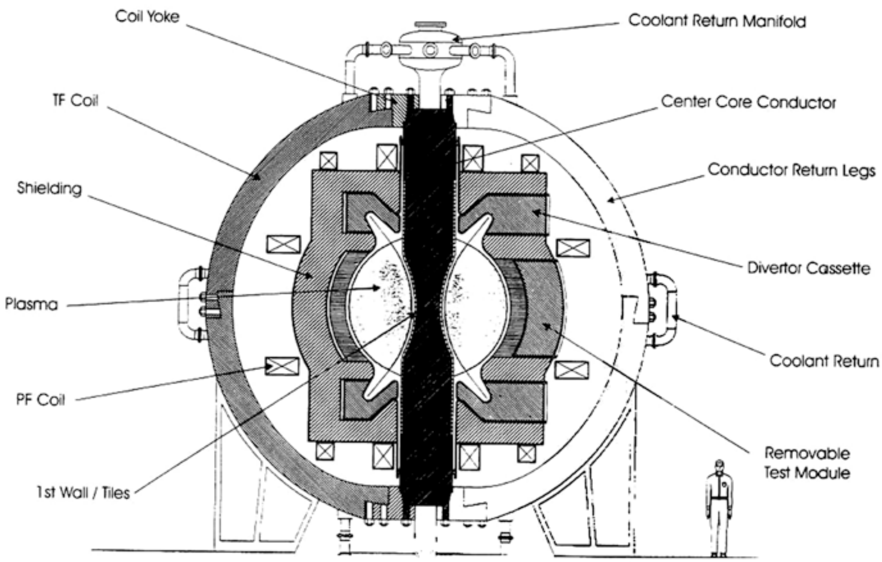

Fig. 9.31 A tokamak neutron source with single-turn normal-conducting toroidal field coils [23]

could cover large areas for blanket development. The group considered both superconducting and normal-conducting toroidal field coils, and it was found that coils made of a single turn rather than multiple windings of copper resulted in a smaller device. This is shown in Fig. 9.31. The major radius is only 80 cm and the toroidal field only 2.4 T; yet the plasma current is 10 MA and the neutron wall loading can be as large as 2 MW/m^2. The last number is indicative of how well the device can duplicate the damage to materials in a reactor like DEMO. This is done well even though the volume neutron source (VNS) is only 0.5% of ITER in volume, 2% in wall area, and 4% in fusion power produced. Significantly, the group did a risk–benefit analysis comparing the ways to obtain an 80% confidence level for DEMO to have, say, 50% availability, taking into account the mean time between failures and the time for repairs. Needless to say, operating ITER with VNS wins hands down over ITER alone. VNS also uses much less tritium in the process. The incremental cost is small: the total of capital cost and operating cost over the life of the machine is $19.6B for ITER and $24.4B for ITER plus VNS.

Fusion Development Facility

A more ambitious tokamak for technology tests has been proposed by a team at General Atomics in San Diego, California [24]. This machine is shown in Fig. 9.32. Note that this depicts only one side of the torus; the major axis is at

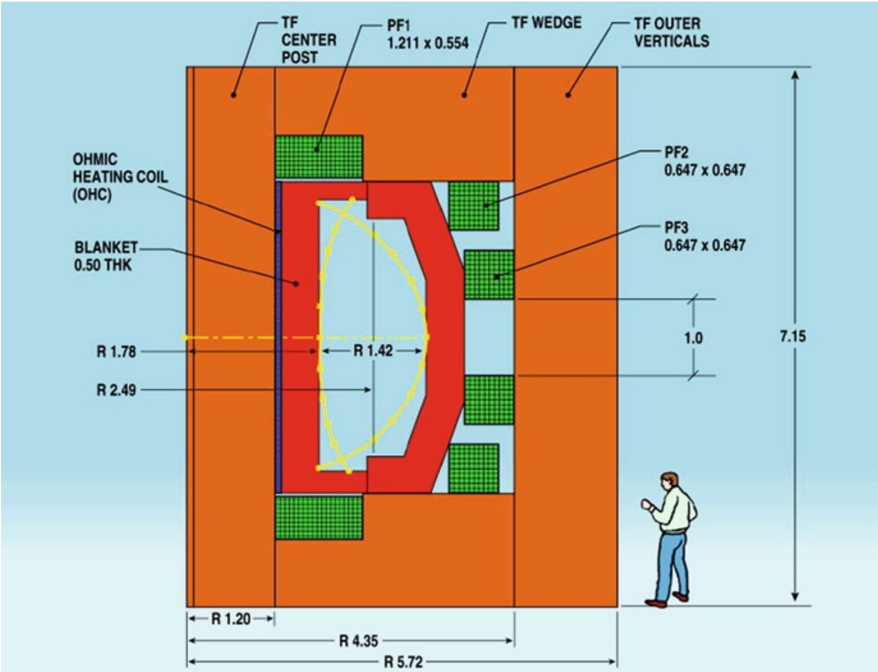

Fig. 9.32 Diagram of the cross section of the FDF tokamak [24]. The centerline of the torus is at the *left* edge of the diagram. TF is toroidal field (coil) and PF is poloidal field (coil). Dimensions are in meters

the left edge of the diagram. The dominant feature is the huge copper toroidal field coil surrounding machine. It will produce a field of 6 T (60,000 G). As seen by the size of the human figure compared to that in Fig. 8.23, FDF is actually smaller than JET. Yet the machine produces 250 MW of fusion power and can run continuously for two weeks at a time. The neutron flux is the required 1–2 MW/m^2, and the fluence is 3–6 MW-years/m^2 over a life of ten years.

Though FDF is much smaller than ITER, it can produce the neutrons for technological testing because it does not reach ignition. It runs steadily at $Q=5$, where Q is the fusion power divided by the power input to the plasma. For ignition $Q>10$ is necessary, and that is much more difficult. Nonetheless, FDF needs all the features of advanced tokamaks: high bootstrap current, internal transport barriers, radiofrequency current drive, and so forth. Remote handling will be developed, with replacement components lowered from the top, where the upper part of the toroidal field coil can be removed. Initially, blanket modules will be tested. Then, after a 2-year shutdown, a full solid ceramic blanket will be installed and tested. In the third stage, after another 2-year shutdown, a Pb-Li blanket will be installed. Only a machine with a full blanket can test such quantities as thermal stress, nuclear waste and disposal, radiation damage, and material lifetimes.

With full blankets, FDF as currently designed can demonstrate a closed fuel cycle, breeding as much tritium as it uses, reaching a TBR of 1.2. In fact, if operated at 400 MW of fusion power, it could actually breed tritium at the rate of 1 kg per year to be stored for use in DEMO. This is a very ambitious goal. In this sense, FDF is comparable to ITER in what it will accomplish. ITER will push superconducting technology, test alpha particle effects, and aim for ignition, but FDF will tackle the harder problems of technology with a smaller machine. FDF will not be cheap at perhaps one-third the cost of ITER; but since it will be a direct replacement for DIII-D, much of the expertise is already in place; and, importantly, the politics of an international project can be avoided. After the cancelation of TFTR, the USA needs to regain its position at the forefront of fusion research.

A Spherical Tokamak FDF

Spherical tokamaks are tokamaks with very small aspect ratio, which is the ratio of major radius to minor radius. They are fat doughnuts with a very small hole in the middle. These are hard to make, but they have advantages in stability. They are described in Chap. 10. Peng et al. [25] have designed a fusion development facility using a spherical tokamak (FDF-ST) with an aspect ratio of 1.5. This is shown in Fig. 9.33. The magnetic coils are normal-conducting copper, even the narrow center leg going through the small central hole. With major radius only 1.2 m, the machine is much smaller than other designs and yet can generate a neutron wall loading of 1.0 or even 2.0 MW/m^2. The toroidal field is 1.2 T, and the plasma current is 8.2 MA. The fusion power is only 7.5 MW or 2.5 times the input power. The machine can accommodate 66 m^2 of blanket area. If this can be engineered, it would be the least costly nuclear test facility to prepare for DEMO.

Fusion Power Plants

Commercial Feasibility

Industry is not interested in these technical details; it is concerned with the bottom line. RAMI is the acronym for four important criteria: reliability, availability, maintainability, and inspectability. The Electric Power Research Institute puts it in even more basic terms: economics, public acceptance, and regulatory simplicity. It is of course too soon to know how these will turn out; but designers of fusion power plants as well as fusion technology researchers are well aware of these criteria, which are always kept in mind. The fusion core is only a part of a whole power plant, a cartoon of which is shown in Fig. 9.34. The remote handling system is essential for maintainability and inspectability. The heating, current drive, and fueling systems affect reliability. The complicated fuel cycle system has to be completely

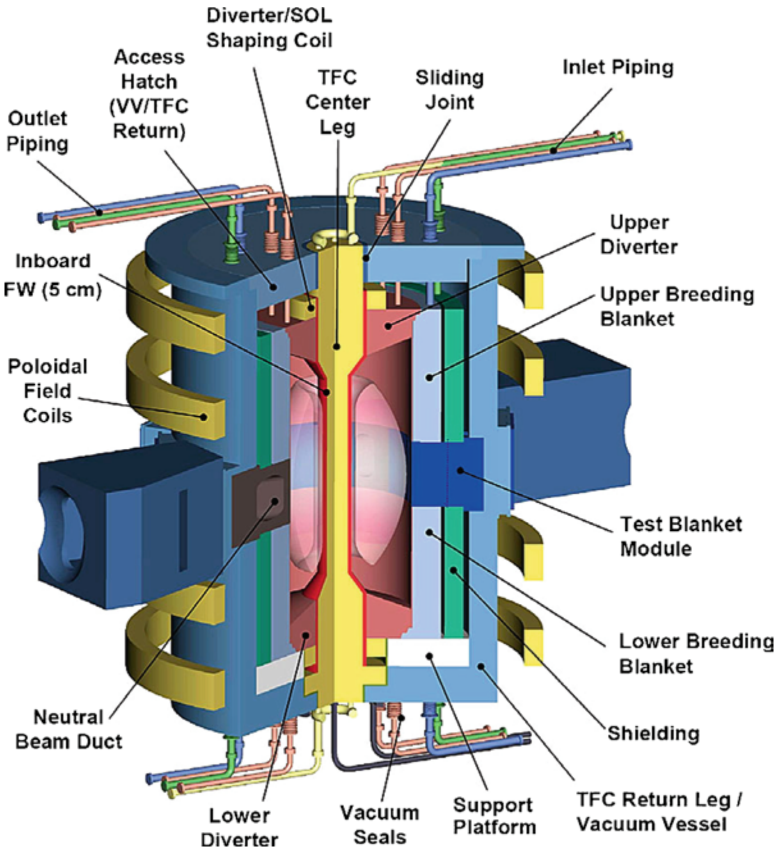

Fig. 9.33 A fusion nuclear science facility using a spherical tokamak [25]

safe in regard to tritium release. The balance of plant, the equipment that generates and transmits the electricity, is a larger part of the power plant than the power core, though it is shown deceptively as a small addition in Fig. 9.34. These are the steam turbines that drive the electric generators and the transformers and capacitors that condition the output for delivery to the transmission lines. All power stations that convert heat into electricity have this equipment, whether the fuel be coal, oil, gas, or uranium. Hydroelectric plants do not need steam; water drives the generators. Wind and solar plants produce electricity directly. Fusion plants can use the same generators and transmission lines that already exist in fossil or nuclear plants; only the power core has to be replaced. However, tokamaks are so complicated and include such temperature extremes that they will require a higher portion of the capital cost than other power cores.

Availability is an important aspect of a fusion reactor that is hard to assess. How often will leaks occur, and how long will it take to do the re-welding? How often do blankets have to be replaced, and how long will the shutdowns be?

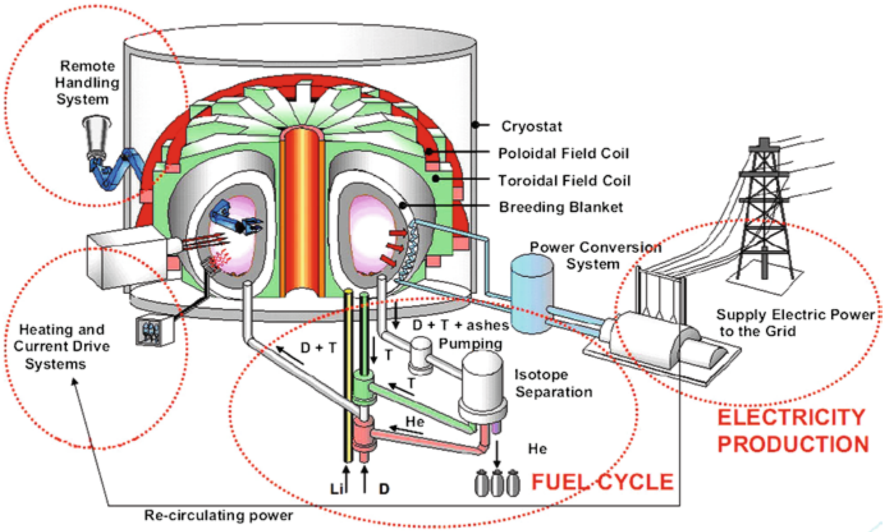

Fig. 9.34 Main parts of a fusion power plant [37]

How often will disruptions occur, and how long will it take to reassemble the machine? What percentage of the time will the machine be running during a year? During a shutdown, where will the power come from? Will we need a backup tokamak or new transmission lines from other power plants? Educated guesses are made by those who design fusion power plants based on available knowledge.

Power Plant Designs

The ARIES program in the USA is the leading group in designing fusion reactors. Originally started by Robert W. Conn in the 1980s at the University of Wisconsin and the University of California (UC) Los Angeles, it is now headed by Farrokh Najmabadi at UC San Diego. Throughout the years, new ARIES designs have been made as new physics has been discovered. The designs are not only for tokamaks; stellarators and laser-fusion reactors have also been covered. The latest designs, ARIES-AT for advanced tokamaks and ARIES-ST for spherical tokamaks, inspired the FDF proposals described above. Practical considerations such as public acceptance, reliability as a power source, and economic competitiveness pervade the studies. The designs are very detailed. They optimize the physics parameters, such as the shape of the plasma and the neutron wall loading. They also optimize the engineering details, such as how to replace blankets and how to join conductors to make the joints more radiation resistant. As new physics and new technology became available, the reactors ARIES I, II, … to ARIES-RS (reversed shear) and

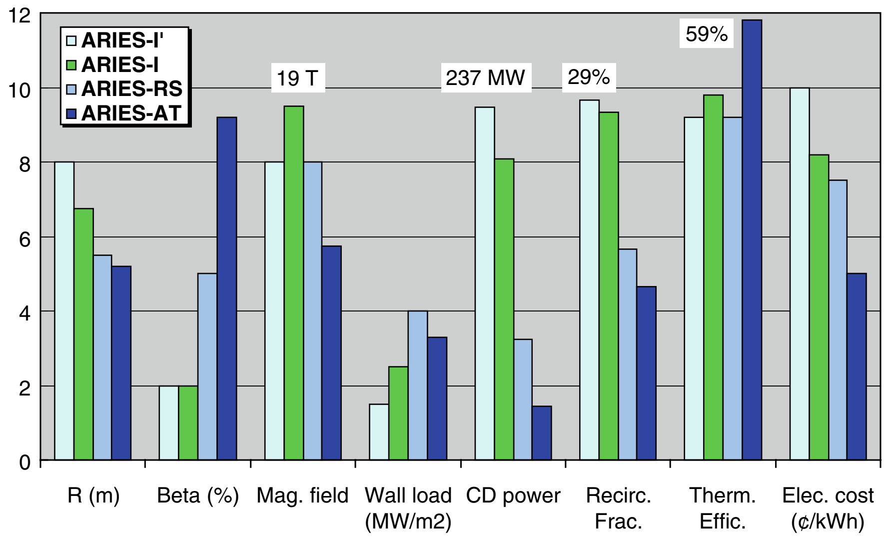

Fig. 9.35 Evolution of ARIES reactor designs. Some bars have been rescaled to fit the chart; these have their maximum values shown. *R* is the major radius. Beta is the ratio of plasma energy to magnetic field energy. The peak magnetic field is given in Tesla. The neutron wall loading is in MW/m^2. CD is current drive. Recirc. Frac. is the recirculating power fraction. Therm. Effic. is the thermal efficiency of the plant. Elec. cost is the cost of electricity in cents per kilowatt-hour

ARIES-AT (advanced tokamak) evolved to become smaller and cheaper. This is shown in Fig. 9.35. We see that as fusion physics advanced from left to right in each group of bars, the size of the tokamak, the magnetic field, and the current-drive power could be decreased while increasing the neutron production. This is due to the great increase in plasma beta that the designers thought would be possible. The recirculating power fraction is the power used to run the power plant; the rest can be sold. It dropped from 29 to 14%. The thermal efficiency in the latest design breaks the 40% Carnot-cycle barrier by the use of a Brayton cycle. Finally, we see that the COE is expected to be halved from 10¢ to 5¢ per kWh with advanced tokamaks.

ARIES-AT is shown in Fig. 9.36. Unlike existing tokamaks, this reactor design has space at the center for remote maintenance and replacement of parts. The philosophy in reactor design is to assume that the physics and technology advancements that are in sight will actually be developed and, on that basis, optimize a reactor that will be acceptable to industry and the public. It is not known whether high-temperature superconductors will be available on a large scale, but this would simplify the reactor. The blankets will be of the DCLL variety, and it is predicted that the Pb-Li can reach 1,100°C without heating the SiC walls above 1,000°C. This high temperature is the key to the high thermal efficiency. For easier maintenance and better availability, the blankets are made in three layers, two of which will last the life

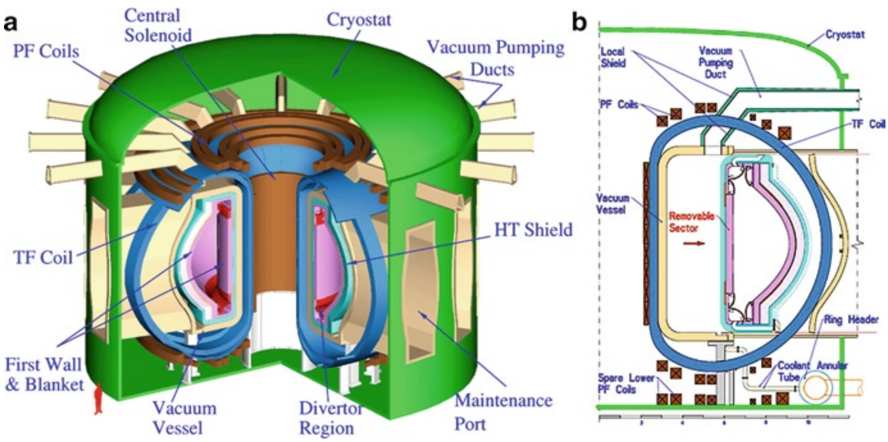

Fig. 9.36 Drawing of the ARIES-AT reactor design and its cutaway view at the *right* [38]

of the reactor. Only the first layer, along with the divertor, has to be changed out every five years. Sectors are removed horizontally and transported by rail in a hot corridor to a hot cell for processing. Shutdowns are estimated to take four weeks.

Turbocharging and supercharging in automobiles are terms that are well known to the public. Airplanes engines are turbocharged. Modern power plants use thermodynamic cycles that have higher efficiency than the classic Carnot cycle. The ARIES-AT reactor will use one of these called a Brayton cycle. The hot helium from the tokamak blanket is passed through a heat exchanger to heat helium that goes to electricity-generating turbines. The two helium loops are isolated from each other because the tokamak helium can contain contaminants like tritium. The turbine also runs with cooler helium at a different flow rate. The Brayton cycle precompresses the helium three times before it goes into helium turbines. The heat of the helium coming out of the turbines is recovered in coolers that cool the helium before it is compressed. It is this system that achieves the 59% thermal efficiency of the ARIES-AT design.

ARIES-AT will produce 1,755 MW of fusion power, 1,897 MW of thermal power, and 1,136 MW of electricity. The radioactive waste generated will be only 30 m^3 per year or 1,270 m^3 after 50 years. The plant will run for 40 of those years if availability is 80%. Ninety percent of this waste is of low-grade radioactivity; the rest needs to be stored for only 100 years. No provisions for public evacuation are necessary, and workers are not exposed to risks higher than in other power plants. The COE from ARIES-AT is compared with other sources in Fig. 9.37. We see that electricity from fusion is not expected to be extravagant.

Europeans have also made reactor models in their Power Plant Conceptual Studies (PPCS) [26]. Figure 9.38 is a diagram of the tokamak in those designs. As with the ARIES studies, Models A, B, C, and D in PPCS (Fig. 9.39) trace the evolution of the design with advances in fusion physics and technology, with Model D using the most speculative assumptions. All these models produce about 1.5 GW of

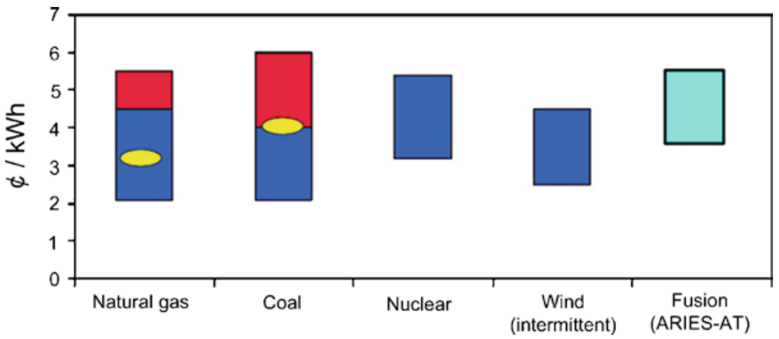

Fig. 9.37 Estimated year 2020 cost of electricity in US cents per kilowatt-hour from different power sources [graph adapted from [25], but original data are from the Snowmass Energy Working group and the US Energy Information Agency (*yellow ellipses*)]. The *red* range is the cost if a $100/ton carbon tax is imposed. The fusion range is for different size reactors; larger ones have lower cost

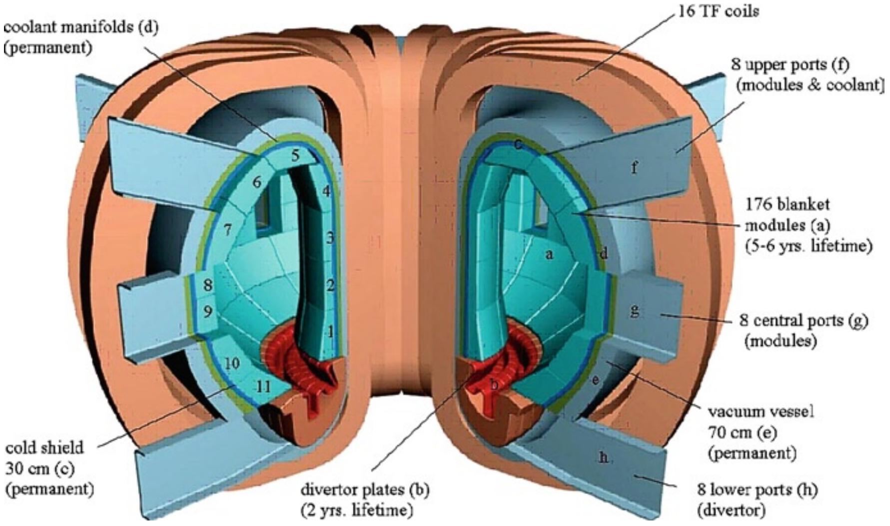

Fig. 9.38 Drawing of tokamak in Power Plant Conceptual Studies in Europe [26]

electricity, but they are smaller and use less power with gains in knowledge. The recirculating fraction and thermal efficiency of Model D matches that of ARIES-AT. Safety and environmental issues were carefully considered. The cost estimates are given in Fig. 9.40, also in US cents per kWh. The difference between the wholesale price of electricity and that available to consumers is clearly shown. It is seen that fusion compares favorably with the most economical sources, wind and hydro.

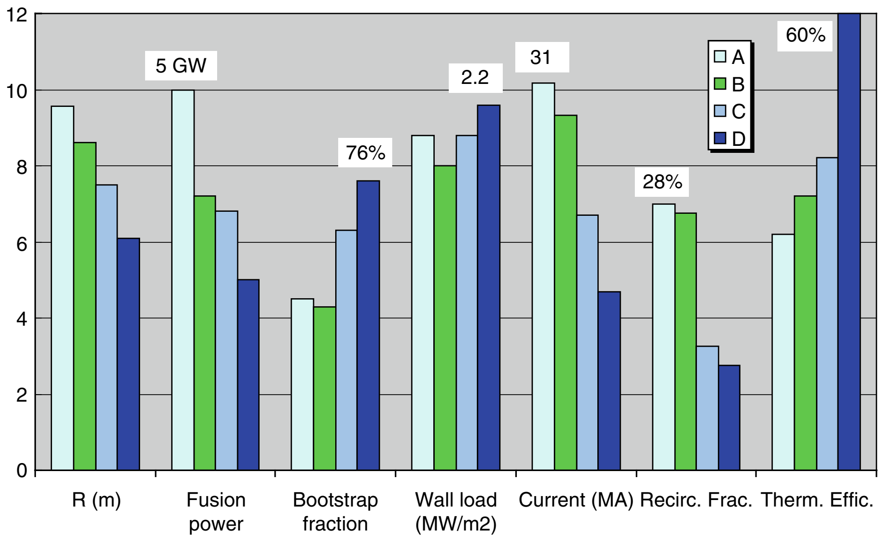

Fig. 9.39 Evolution of PPCS designs [26]. See caption of Fig. 9.35

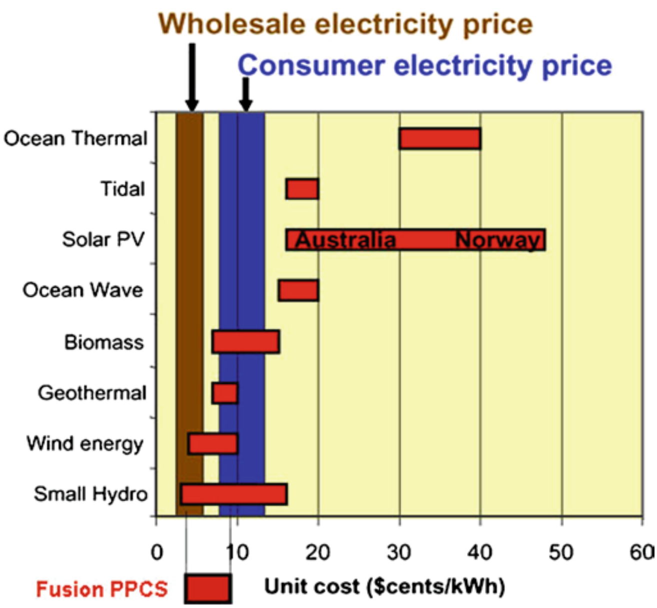

Fig. 9.40 Cost of electricity from fusion compared with other renewable sources [26]

The Cost of Electricity

Methodology

In spite of the fact that we do not yet know how a fusion reactor will be constructed, or even if it is at all possible, detailed calculations have been made on the COE based on the reactor models described in the previous section. The work of Ward et al. [27], which we will summarize here, is based on the European PPCS designs. Their calculated costs for each component of a power plant compare well with those from the ARIES studies in the USA. Being a renewable power source, fusion shares with wind, solar, and hydro the benefit of essentially zero fuel cost. However, the capital cost is large. A breakdown is given by Ward [28] in Fig. 9.41. The capital cost of the tokamak power core is almost as large as that of the balance-of-plant, which is the power conversion system and electrical generators shown in Fig. 9.34. Compared with fossil fuel plants, the capital cost and replacement of blankets and divertors take the place of fuel costs. These fusion-specific costs depend on the reactor model. The models A, B, C, and D in Fig. 9.39 range from ITER-like primitive designs with steel chambers and water cooling to speculative advanced designs with Pb-Li liquid cooling and SiC/SiC first walls. Computer programs are used to calculate the costs of each component under different assumptions.

Important Dependences

The COE depends on some factors that are independent of the power core and others that are specific to fusion. These factors appear in the following formula for COE,

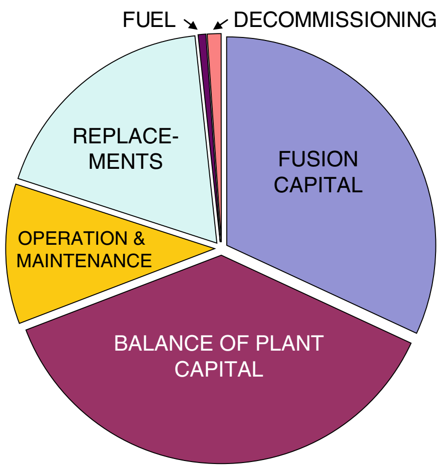

Fig. 9.41 Cost breakdown of a fusion power plant [28]

which at first seems rather daunting. However, it is not necessary to know what the formula means in detail; it is used here just as a convenient way to show what affects the cost. The COE is proportional to the quantities in the parenthesis times those in the

$$\mathrm{COE} \propto \left(\frac{rL}{A}\right)^{0.6} \frac{1}{\eta_{\mathrm{th}}^{0.5} P_{\mathrm{e}}^{0.4} \beta_{\mathrm{N}}^{0.4} N^{0.3}}$$

denominator of the fraction following it. Inside the parenthesis, r is the discount rate, a financial factor similar to interest rate that will be explained later. L is a learning factor which takes into account that the first of a kind is always more expensive than the tenth one made. L starts at 1 and gets smaller with learning, so COE drops. A is the availability, which is the fraction of time the plant is running rather than shut down for repairs. Larger A means lower costs. The fusion reactor designs have A's ranging from 60 to 80%.

The first two quantities in the denominator at the right have to do with the whole plant, and the last two concern the quality of the plasma in the tokamak. Eta-thermal (η_{th}) is the efficiency of converting heat into electricity. P_{e} is the size of the plant in terms of electrical power produced. The larger the better because of economy of scale. Beta-normalized (β_{N}) expresses the efficiency with which the plasma current can confine a large amount of hot plasma by creating the right amount of twist in the magnetic field. Finally, N is the ratio of the plasma density to that predicted by Greenwald limit (Chap. 8) for a stable plasma. In the different reactor models, r varies from 5 to 10%, L from 0.5 to 0.7, A from 0.6 to 0.8, η_{th} from 35 to 60%, P_{e} from 1 to 2.5 GW, and N from 0.7 (safe) to 1.4 (speculative). Most importantly, β_{N} varies from 2.5 to 5.5, representing the progression from well-established data to hopefully achievable advanced tokamak operation. Figure 9.42 shows the COE predicted from the PPCS models A–D as a function of the learning factor L.

As an example of how sensitive the COE is to assumptions made in the models, Fig. 9.43 shows how the availability factor A changes with the lifetime of the materials

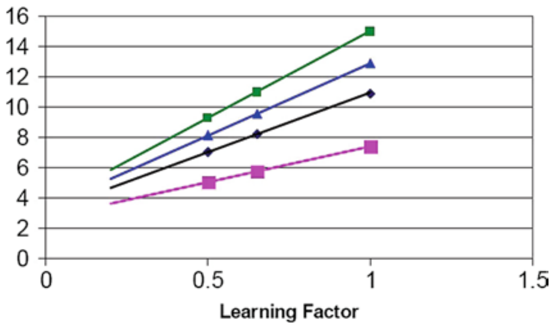

Fig. 9.42 The cost of electricity, in euro cents per kilowatt-hour, calculated for various reactor models as a function of the learning factor L [28]. Model A is an ITER-like machine, and D is the most advanced reactor envisioned at present. Power plants start at $L = 1$ and progress leftward to lower costs as more are built

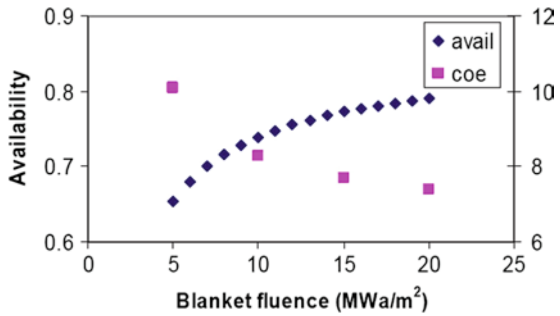

Fig. 9.43 Dependence of the power plant availability and cost of electricity on the degree to which materials in a tokamak reactor can withstand neutron damage [28]

in the first wall and blankets. The lifetime is expressed as the neutron fluence that the materials stand before they have to be replaced. The fluence is measured in years at an equivalent neutron energy flux of 1 MW/m^2. The shorter the lifetime, the more often the blankets will have to be replaced, and hence the lower the availability. This then increases the cost (the higher blue points at the left).

Cost Levelization/Discounting

Expenses and income are both functions of time. Costs start accruing when a power plant is proposed and initial studies are made, for instance, on environmental impact. Land is purchased, the plant is designed, equipment is ordered, and construction begins. This takes many years. The plant is finished and begins producing power. Profits begin to be made, year by year. At the same time, there are expenses for operating the plant, and for repairing and replacing equipment. To get a reasonable number for the COE, one has to adjust all the expenses and income forward or backward to the same date. Time is money. This is called discounting. It is done with another formula:

$$\mathrm{COE} = \frac{\Sigma_t (C + \mathrm{OM} + F + R + D)_t / (1+r)^t}{\Sigma_t E_t / (1+r)^t}$$

This is a formula unfamiliar to physicists but may be more familiar to readers involved with business or finance. Here C is the capital cost, OM is operation and maintenance, F is the fuel cost, R is the cost of replacements, D is the cost of decommissioning at the end of life, and r is the *discount rate*. In the denominator, E is for earnings. The sum is over time t. To derive a value at time zero for an expense or income occurring at another time, a discount has to be applied. The discount rate is like an interest rate but includes also expectations of what the market will be like, how much inflation there will be, and factors like those. Financiers normally assign a discount rate between 5 and 10%.

Suppose we want to calculate the COE as of the start of planning. We set that as $t=0$. For simplicity, let us do the accounting annually, not monthly or daily. Suppose it takes five years to get ready, five years to build the plant, and it has been operating for another five years. For years $t=1$–5, we have the money C_1-C_5 spent in those years, which is only interest on money borrowed, salaries, and rental for office space. For years $t=6$–10, C will be much larger, as the plant is built. For years 11–15, we have $C+\mathrm{OM}+F$ for those years, and also E earned in those years. Each year's amounts are divided by $(1+r)$ raised to the power t in order to get the value as of $t=0$. Both the numerator and the denominator are summed over the years, and the ratio is the COE. In later years, there will also be values for R and D.

To get a better idea of what discounting means, let us consider a simple example. Suppose you borrow \$1M to build a machine, taking five years to do so. At the end of the five years you sell the machine for \$1M. However, you could not have sold that machine for \$1M at Year 0, since that machine did not exist yet and you could not make any money with it. It has a smaller discounted value at Year 0, given by $C/(1+r)^5$, according to the formula above. If $C=\$1\mathrm{M}$ and the discount rate is $r=5\%$, we have a value at $t=0$ of $C/(1.05)^5$, which works out to be only \$0.784M. The reason is that you had to pay compound interest during the five years. One million dollars compounded annually at 5% is \$1M times $(1.05)^5$, which is \$1.276M. You had to pay \$0.276M in interest, so you made only \$0.724M, and that is closer to the value of the machine at $t=0$. Actually, you did not have to borrow all the money at once, so the discounted, or levelized, value is \$0.784M, which is exactly the reciprocal of \$1.276M.

This exercise points out that a large part of the cost of any power plant, regardless of its power source, is interest during construction. If the discount rate is 7.5% (halfway between 5 and 10%), and the plant takes five years to construct, summing over the discounted value of one-fifth of the capital cost for each of five years shows that 20% of the cost is interest and other financial factors. The levelized COEs of all different kinds of power plants (except fusion) in many different countries have been analyzed in exhausting detail by the International Energy Agency.[6]

The Cost of Fusion Energy

Figure 9.44 shows how the COE from fusion compares with that from other energy sources [28]. Each entry has two bars showing a minimum and a maximum value, the difference depending partly on location and partly on technology. For fossil fuels, the maximum is the cost including the expense of carbon sequestration. For fusion, the maximum and minimum represent the range of the reactor models ABCD described above. These data for other energy sources are from the IEA report of 1998. Fuel prices and interest rates have fluctuated so violently in recent years that the comparison has not been updated. However, the levelized COEs of nonfusion sources are available for 2005[6] and 2010.[7] The data for 2010 are shown in Fig. 9.44. For comparison, the fusion COE given in Fig. 9.44 is reproduced in Fig. 9.45. That graph shows also the breakdown between capital costs and

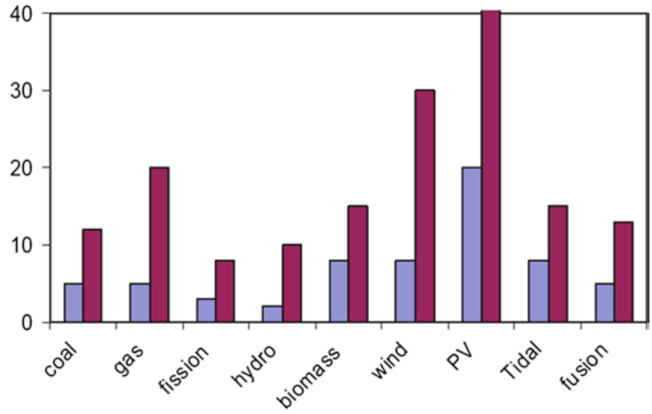

Fig. 9.44 Comparison of the cost of electricity from conventional and renewable energy sources [28]

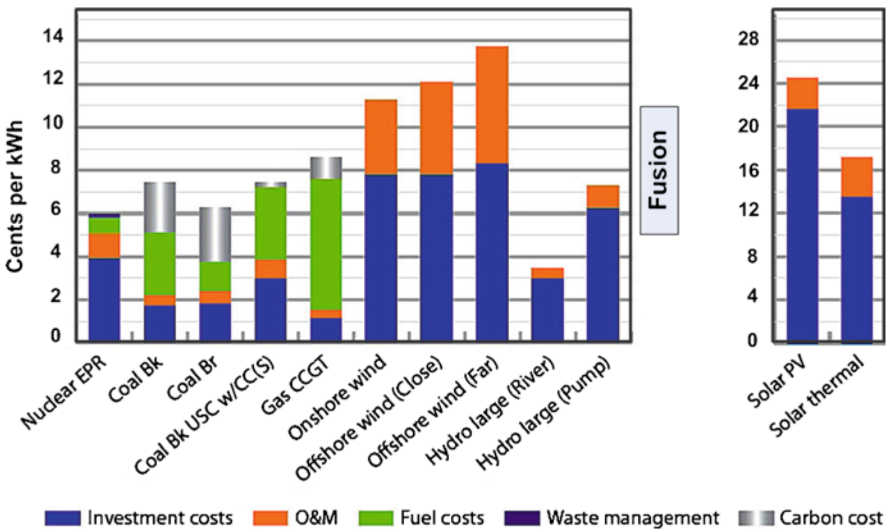

Fig. 9.45 Estimated cost of electricity in Europe from nuclear, fossil-fuel, and renewable sources assuming a 5% discount rate[7]. The color code gives the breakdown among capital costs, operation and maintenance (O&M), and fuel costs. For nuclear plants, there is charge for nuclear waste management. For fossil-fuel plants, there is a cost for carbon management under certain assumptions. The estimated cost range for fusion plants has been added. The solar photovoltaic (PV) and solar thermal costs have to be plotted on a different scale

operation and maintenance costs, as well as the estimated cost of carbon capture and sequestration for fossil-fuel plants. The data are from different time periods, but the difference is insignificant in view of the uncertainties involved. It is seen that *the COE from fusion plants will be competitive with that from other renewal sources and from fossil-fuel plants with carbon management.*

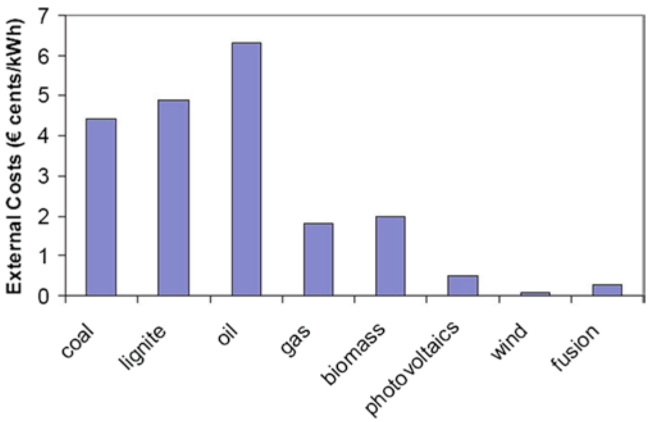

Fig. 9.46 External costs of fusion compared with other energy sources [27]

It is interesting to note that the large variability of the COE is reflected in the IEA's 2010 report[7]. The figures for each energy source vary greatly from country to country. In addition, the sensitivity of the estimates to such factors as corporate taxes, discount rate, and fuel cost is emphasized.

Not included in the above analyses are external costs, which include damage to the environment, general health, and human life. Such costs have been evaluated by site to eliminate location biases. For instance, one considers the difference when a fusion plant is put in place of a coal plant in the same location. It turns out that the external costs of fusion are extremely low, ranging from 0.07 to 0.09 euro cents per kWh. Comparison with other energy sources is shown in Fig. 9.46.

The net present value of fusion takes into account the probability of success or failure. Though this obviously has a high degree of uncertainty, there is a large margin for error, since the annual world energy expenditures exceed the annual cost of fusion development by 1,000 times. It has been estimated that if fusion captures 10–20% of the electricity market in 50 years, the discounted future benefit of fusion is $400–800B; or, if the probability of failure is counted, it is still $100–400B. This means that development of fusion is worthwhile even if fusion captures only 1% of the world electricity market [27].

Notes

1. However, a vertical Allure Ignition Stellarator with a liquid Li wall was proposed in 2010 to be built in Spain.
2. It has been pointed out that tritium is also generated in the beryllium multiplier, an effect usually neglected in estimates of breeding ratio [3].
3. ITER Newsline Nos. 114 and 122 (2010). http://www.iter.org/newsline/.
4. M. J. Schaffer (General Atomic), private communication.
5. A. Möslang (Karlsruhe), *Strategy of Fusion Materials Development and the Intense Neutron Source IFMIF*.

6. *Projected costs of generating electricity, 2005 update*, published by the Nuclear Energy Agency and the International Energy Agency of the Organization for Economic Cooperation and Development (OECD).
7. *Projected costs of generating electricity, 2010 Edition*, published by the Nuclear Energy Agency and the International Energy Agency of the Organization for Economic Cooperation and Development (OECD).

References

1. L.L. Snead, *Ceramic Structural Composites, the Most Advanced Structural Material, 9th Course on Technology of Fusion Tokamak Reactors* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
2. L. Giancarli et al., Fusion Eng. Des. **61–62**, 307 (2002)
3. L. Giancarli, *The PPCS In-Vessel Component Concepts* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
4. D. Stork, *DEMO and the route to fusion power*. 3rd Karlsruhe International School on Fusion Technology, Sept 2009
5. A. Ying et al., *Current status of the test blanket program in ITER and implications for blanket component testing requirements and goals*. Symposium on Fusion Energy, San Diego, CA, June 2009
6. G. Janeschitz, *The development of commercial fusion energy in the EU*. Seminar, University of California, Los Angeles, Jan 2008
7. S. Casadio, *Ceramic Breeder Technology* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
8. A. Abou-Sena, A. Ying, M. Abdou, Fusion Sci. Technol. **56**, 206 (2009)
9. J. Takeuchi et al., Fusion Eng. Des. **83**, 1082 (2008)
10. S. Smolentsev, R. Moreau, M. Abdou, Fusion Eng. Des. **83**, 771 (2008)
11. A. Pizzuto, *Comparison of Breeder Blanket Designs* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
12. M. Glugla et al., *Review of the ITER fuel cycle systems and recent progress*. Symposium on Fusion Energy, San Diego, CA, June 2009
13. I.R. Cristescu, *The Fuel Cycle System* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
14. E. Salpietro, *Magnet Design and Technology* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
15. Report of the Research Needs Workshop (ReNeW), Bethesda, MD, June 8–12, Office of Fusion Energy Sciences, US Department of Energy, 2009
16. P.J. Heitzenroeder et al., *An overview of the ITER in-vessel coil systems*. Princeton Plasma Physics Laboratory Report PPPL-4465, Sept 2009
17. M.J. Schaffer et al., Nucl. Fusion **48**, 024004 (2008)
18. L.R. Baylor et al., *Disruption mitigation technology concepts and implications for ITER*. Symposium on Fusion Energy, San Diego, CA, June 2009
19. H.L. Berk et al., Phys. Rev. Lett. **68**, 3563 (1992)
20. G. Giruzzi et al., Nucl. Fusion **49**, 104010 (2009)
21. J. Bucalossi et al., *Performance issues for actuators and internal components during long pulse operation*. Symposium on Fusion Energy, San Diego, CA, June 2009
22. J. Rathke, *Engineering overview: International fusion materials irradiation facility*. FESAC Subcommittee Review (FESAC is the Fusion Energy Sciences Advisory Committee of the US Department of Energy), San Diego, CA, Jan 2003
23. M.A. Abdou et al., Fusion Technol. **29**, 47 (1996)
24. A.M. Garofalo et al., IEEE Trans. Plasma Sci. **38**, 461 (2010)

25. Y.K.M. Peng et al., Fusion Sci. Technol. **56**, 957 (2009)
26. D. Maisonnier, *PPCS Reactor Models* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
27. D.J. Ward et al., Fusion Eng. Des. **74–79**, 1221 (2005)
28. D.J. Ward (EURATOM/IKAEA Fusion), *Impact of Physics on Power Plant Design and Economics* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
29. N. Holtkamp, *Status of ITER*. Fusion Power Associates meeting, Washington, DC, Dec 2009
30. G. Federici, *Plasma Wall Interactions in ITER and Implications for Fusion Reactors* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
31. G. Janeschitz, *Divertor Physics and Technology* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
32. M. Abdou, *Challenges and development pathways for fusion nuclear science and technology*. Seminar, Seoul National University, S. Korea, Nov 2009
33. N.B. Morley, M. Abdou, in *Fusion Power Associates Annual Meeting*, Washington, DC, Oct 2005
34. D. Ward (UKAEA, Culham, UK), *Impact of Physics on Power Plant Design and Economics* (International School on Fusion Reactor Technology, Erice, Italy, 2004)
35. K. Sakamoto et al., Nucl. Fusion **43**, 729 (2003)
36. G.T. Hoang et al., Nucl. Fusion **49**, 075001 (2009)
37. J. Pamela et al., *Key R&D issues for DEMO*. Symposium on Fusion Technology, Rostock, Germany, Sept 2008
38. F. Najmabadi, in 18th KAIF/KNS Workshop, Seoul, Korea, 21 April 2006