FUSION REACTORS

NUCLEAR REACTORS, FUSION

REACTORS, FUSION

NUCLEAR FUSION

The fission plants which now provide about 20% of the world's electricity accomplish this feat through neutron-driven chain reactions in which heavy atomic nuclei, such as U^{235}, split apart into lighter nuclei, releasing large amounts of energy (on the order of 180 million eV) in the process. Beginning about 1951, when one of the first fusion programs was begun at Princeton University by Lyman Spitzer, Jr., many of the world's developed nations, as well as some developing ones, have pursued research to eventually produce a fundamentally different type of nuclear reactor. Since 1958, when research on peaceful uses of nuclear energy was declassified, these many countries have shared their research and openly collaborated, even during decades when some of the principal contributing nations were political adversaries. This new type of reactor is a fusion reactor, and the very great progress which has been made in the decades since 1951 stands as a landmark to what may be the most universal and long-lived collaborative effort in human history.

Fusion reactors, like fission reactors, will use exothermic nuclear reactions to release energy. The fusion reactions themselves, however, are very different in character than are those in fission reactors, and they require entirely different conditions in order to proceed at an acceptably high rate for power plant applications. Due to these differences, fusion reactor designs will look very different from those of fission reactors. Moreover, many of the problems that have had to be surmounted in the pursuit of fusion power are of a fundamentally different nature than those encountered in making fission reactors practical.

Nuclear fusion is a reaction in which two atomic nuclei merge to form a heavier element. If the reaction is an exothermic one, then the fusion process will result in the release of energy. This energy is carried as kinetic energy by the reaction products, consisting of the product nucleus (which is positively charged) and another particle, such as a neutron (which is electrically neutral) or a proton (which is positively charged). Because all nuclei carry positive electric charge, they repel each other, with the result that fusion reactions can proceed at significant rates only at very high temperatures which give the nuclei sufficient energy to overcome their repulsion and approach each other close enough to merge. Alternatively, this same end may be achievable by using somewhat lower temperatures, but at high densities achieved by compressing the fusion fuel with high symmetrically applied pressures.

The electrostatic repulsive force between two nuclei is proportional to the product of the positive charges (and thus to the atomic numbers) of the two reactant nuclei. This gives rise to a potential barrier, referred to as a coulomb barrier, which the approaching nuclei must overcome in order to merge. Due to quantum mechanical tunneling, some pairs of nuclei can fuse even when their combined kinetic energy is less than that required to exceed the coulomb barrier. However, the likelihood of this tunneling occurring declines very rapidly as the kinetic energy of the reactants falls farther below the Coulomb barrier height. For most nuclear reactants, the reaction rate drops to inconsequential levels if the approaching nuclei do not have kinetic energies of at least 70% to 80% of the barrier height. For heavier nuclei, the height of the Coulomb barrier becomes several to many millions of electron volts (*MeV*), where an electron volt is a unit of energy corresponding to the energy one electron gains when it falls through a potential difference of 1 V. A temperature of 1000 eV (1 keV) corresponds, in more conventional units, to about 10,600,000°C.

Thus, in order to fuse most of the elements in the periodic table, one would need to cause them to approach each other at energies of several to many MeV. This can be accomplished with particle accelerators for nuclear physics research purposes, but is not practical for producing net power with significant quantities of thermal reactants. As a consequence, only the lightest elements in the periodic table have sufficiently large probabilities (also called reaction cross sections) of undergoing nuclear fusion for them to be considered as fuels for fusion reactors.

Fusion reactions drive the core of the sun at a temperature thought to be roughly 10 million to 11 million °C. Fusion also powers all the other stars, and thus supplies most of the universe's observable energy. Our sun, in common with other main sequence stars, obtains its energy through a number of nuclear fusion reactions, beginning with the fusion of 2 protons into deuterium. The deuteron is accompanied by a positron and a neutrino, and together they carry 1.44 MeV of energy. The reaction probabilities (or cross sections) for this and the succeeding solar nuclear reactions are much too small to be of any use to a commercial fusion reactor on the earth's surface. The sun is able to produce enormous amounts of energy with these reactions only because it is very large compared to the earth. The sun has a diameter of about 897,000 mi and, because of its size and its gravity, has excellent energy confinement. The time for energy produced in the sun's fusion driven core to reach the sun's surface is estimated to be of the order of thousands of years or more.

Because a practical fusion reactor for electricity production needs to be much smaller than the sun, it will have to rely upon different nuclear fusion reactions with larger cross sections, and it will also need to operate at temperatures that are at least 10 to 20 times those of the sun's core. Temperatures in this range were achieved within a mirror confinement device in 1979 and within the Tokamak Fusion Test Reactor (*TFTR*) at Princeton University in the mid-1980s, and have now also been obtained on the Joint European Torus (*JET*) tokamak in England, and the JT-60U tokamak of Japan. Ion temperatures as high as 300 million to 500 million °C (roughly 30 to 50 times the temperature of the sun's core) were routinely produced in Princeton University's TFTR over a period of years.

FUSION REACTOR FUELS

There are a number of nuclear fusion reactions which have cross sections sufficiently large to be potential candidates as fuels for commercial fusion reactors. These are listed below, where D represents a deuteron, the nucleus of ^{2}H, T stands for a triton, the nucleus of ^{3}H, and the energies in parentheses are the amounts of kinetic energy each of the reaction products carries away from a fusion event. An n stands for a neutron, and a p for a proton, the two types of baryons which occur in atomic nuclei.

$$\mathrm{D + T \rightarrow {}^4He(3.52\,MeV) + n(14.06\,MeV)}$$
$$\mathrm{D + D \rightarrow \ T(1.01\,MeV) + p(3.02\,MeV)}$$
$$\mathrm{D + D \rightarrow {}^3He(0.82\,MeV) + n(2.45\,MeV)}$$
$$\mathrm{D + {}^3He \rightarrow {}^4He(3.67\,MeV) + p(14.67\,MeV)}$$

Of these reactions, the D + T reaction has by far the largest cross section at energies of tens of keV, which is the region which should be obtainable in the first generation of commercial fusion reactors. For a fusion reactor with a thermal reacting population, the parameter which describes the fusion reaction rate is $\langle \sigma v \rangle$, where σ is the nuclear reaction cross section, and v is the relative speed of the reacting nuclei. The angle brackets mean that the product of σv is averaged over the Maxwellian velocity distributions of the reactants. This fusion reactivity parameter is also much larger for the D + T reaction than for the other possible fuel mixtures. Due to the importance of the reactions arising from the high energy tail of the Maxwellian velocity distributions, the optimum range for ignition of a D + T fuel mix lies across a range of about 20 keV to 40 keV, well within the operating range of some of the largest of existing fusion devices. These experimental devices have not actually reached ignition because, although their temperatures were adequate, the product of the density and the energy confinement time was not yet sufficiently large. The energy confinement time is a measure of how long is required for energy to leak from the plasma. In the 1950s a typical experimental device had an energy confinement time of a few milliseconds for the overall plasma; by the mid-1990s this reached values of as much as 1.4 s.

In addition to the magnitude of the reactivity as a function of temperature, there are also other factors which bear upon the ease with which different fusion fuels can be used in a commercially viable reactor. There are two types of radiation loss which take energy from the confined fuel and transport it directly to the outside, and both of these tend to be more serious for higher temperatures. The first of these is bremsstrahlung, which scales in intensity proportional to the square of the density times the square root of the temperature. For D + T fueled reactors, this would be a relatively tolerable portion of the overall power balance. However, for the other fuels it would play a larger role. For approaches to fusion reactors which employ strong magnetic fields, the electron synchrotron radiation becomes a major factor at the higher temperatures (60 keV to 120 keV) required by reactions involving D + D or D + ^{3}He. Thus, the appeal of a D + T fuel mix in at least the early electricity-producing fusion reactors arises from two major factors: the fusion reactivity is much higher at temperatures which can be reached by some of today's fusion devices, and the radiation loss rates from bremsstrahlung and synchrotron radiation are lower for likely reactor conditions. Synchrotron radiation would be less important for alternate plasma confinement schemes which do not use strong magnetic fields, but these approaches have not been developed as far as the strong field path to fusion.

Deuterium Fuel

Deuterium is a stable, naturally occurring isotope of hydrogen. On the earth, one out of every 6700 atoms of hydrogen is deuterium. Thus, enormous supplies of deuterium are available from the earth's water. If later generations of reactors operate with the D + D reaction, then there is enough fuel to supply the world's energy needs for hundreds of millions to billions of years, depending upon assumptions about the future growth of energy usage. The deuterium can be concentrated and extracted from any water by utilizing such enrichment techniques as diffusion through a series of filters or by electrolysis.

The earth's water contains about 4×10^{16} kg of deuterium. If this were used to fuel D + T fusion reactors with an overall operating efficiency of 33%, then this would allow the production of about 10^{22} *GJ* of electricity, which is close to 3×10^{11} the present annual electricity production of the entire world. As we will see later in this section, the real limit on the amount of energy potentially available from D + T fusion reactors is the supply of feedstock to produce the tritium used in the reaction.

If later reactors used D + D as their primary reaction, then, because this reaction is less exothermic that the D + T reaction, and because two deuterium atoms, instead of one, would be required for each reaction, then the estimate of total electricity production if all the deuterium in the waters of the world were used would drop to about 10^{21} GJ, which is still a very large number, and is equivalent to about 3×10^{10} times the world's current annual production of electricity. The cost of deuterium is of the rough order of one dollar per gram (1), with a gram of deuterium being sufficient to produce 300 GJ of electricity if fusioned with tritium. The cost of the deuterium fuel for a D + T reactor is thus about \$0.003 per gigajoule of electricity. If the deuterium were instead used in a D + D reactor, the cost would rise to about \$0.02 per gigajoule of electricity, which is still small compared to the cost of bulk electricity, which runs in the vicinity of \$20 per gigajoule. The fuel costs would be a negligible portion of the cost of electricity from a fusion plant, which would be dominated by capital costs and maintenance. It is likely that the price of electricity from a fusion reactor would be similar to or perhaps somewhat more expensive than electricity from a fission reactor, at least in the near term, while there is still adequate fuel for the simple once-through fuel cycles used in most fission power plants.

Comparison to Fossil Fuel Energy Densities

Producing a gigawatt (10^9 W) of electricity for a year in a fusion reactor would require roughly 1000 kg of deuterium. Producing the same amount of electrical energy from a power plant burning coal would require about 2×10^9 kg of carbon. The actual weight of coal required would of course be greater than this, since coal contains other elements besides carbon. The fact that the mass of fuel which has to be carried from the fuel concentration source to a fusion power plant is more than a factor of a million smaller than for competing fossil fuel plants is a significant advantage. It means that moving the fuel for an entire fusion economy would impose no requirements upon the transportation infrastructure, since the masses being moved would be thousands of kilograms instead of billions of kilograms.

The fact that fusion reactors will be able to produce a gigawatt-year of electricity while using fuel which weighs more than a million times less than that required for a fossil fuel burning plant reflects the large difference between nuclear binding energies and electron binding energies. When a reaction occurs between two parent particles, the extra kinetic energy carried by the daughter particles comes from the change in the overall binding energy. This kinetic energy is distributed among the daughters in accordance with the laws of energy and momentum conservation, with the result that the lighter daughter carries most of the kinetic energy.

Burning a fossil fuel is a chemical reaction, which involves exchanges and rearrangements of the electrons outside the nuclei. The outer electrons of atoms typically have binding energies of several to a few tens of electron volts. Thus, the changes in the net binding energy that occur in chemical reactions, such as burning fossil fuels, are typically only a few electron volts. Nuclear binding energies, on the other hand, are typically many millions of electron volts. This has the consequence that the changes in binding energy involved in nuclear reactions are also much greater than for chemical reactions. For instance, in the D + T reaction, 17.58 MeV of energy is released as kinetic energy in inverse proportion to the masses of the daughters. Thus, because a neutron weighs about one fourth as much as does an ^{4}He nucleus, the neutron carries four times as much of the kinetic energy.

The disparity in binding energies between nuclear and chemical reactions arises in turn from the fact that the strong nuclear force is much more powerful across nuclear dimensions than is the coulomb force (which binds the negatively charged electrons to the positively charged nucleus) across dimensions typical of atoms. Consequently, the huge difference in the magnitudes of potential energy available per unit of mass available from nuclear, as opposed to chemical reactions, arises from a major difference in the strengths of two of the fundamental forces of nature. Thus, there is no possibility that any chemical reaction involving changes in electron configurations could ever begin to approach the energy releases available from nuclear reactions.

Tritium Fuel

Tritium, the heaviest isotope of hydrogen, is unstable. It decays by emitting a beta particle (an electron) with an average energy of 5.7 keV. This beta particle is easily stopped by even a piece of paper, so it does not pose a hazard unless it is ingested. Even then, unless large amounts (millicuries or more) are taken into the body, it is not very likely to produce ill effects. This is the result of two factors. The first is that the ionizing radiation released into the body by each tritium decay is much less than the decay energies of fission products (which are typically at least hundreds of keV) or the energies in the decay chains of heavy elements, which can run to over 10 MeV). The second factor is that most of any tritium absorbed by the body would enter as water, and water is continuously excreted, with a biological half-life in human bodies of about 12 to 13 days. The absorption rate for tritium breathed into the lungs as gaseous molecular hydrogen is very low, roughly 20,000 to 25,000 times less than for tritium in water molecules.

Thus, while tritium should always be treated with due care, its possible health effects if mishandled do not appear to be significantly worse than those of many other chemicals routinely handled by an industrialized society. Indeed, tritium is already used in conjunction with phosphors to provide light without the need of electricity in several applications such as school exit lights, some airfield landing lights, and some modern illuminated watches. These applications incorporate significant amounts of tritium. Tritium-powered school exit lights typically use about 15 to 25 curies of tritium, and emergency runway landing lights use much more. This compares with an amount on the order of 100 to 150 curies injected into the Tokamak Fusion Test Reactor at Princeton University for a high-power fusion shot. This device has operated on a routine daily basis using deuterium and tritium for experiments over a period spanning three and a half years without any significant incidents involving tritium contamination of personnel or the environment. This demonstrates that tritium can be handled on a large fusion system without unduly impeding the manner in which it is operated, although the total quantities of tritium being handled were orders of magnitude lower than would be required in a power plant. The neutron-induced activation and damage of materials would be much more challenging in a fusion power plant.

While the beta decay of tritium does not result in either undue hazard potential or in excessive constraints upon the operation of fusion reactors, it does have an inconvenient consequence. The half-life (the time for half of the nuclei in any assemblage to undergo beta decay) is only 12.3 years. Thus, any primordial tritium that fell into the earth during its formation decayed away billions of years ago. Some tritium is continuously produced in the upper atmosphere through nuclear reactions initiated by cosmic rays. However, due to the short half-life of tritium it does not build up, so the equilibrium concentration of tritium in air is very low, and far too small to economically extract as fuel for a fusion fuel.

Consequently, tritium fuel for fusion must be manufactured through nuclear reactions. The tritium used for present-day fusion experiments is made in heavy water nuclear fission reactors. However, the amount produced by these reactors would be insufficient to fuel a fusion economy. Equally important, the long-term goal of fusion research is to produce fusion reactors which supplant fission reactors as we know them, so fusion reactors must produce their own tritium. Fortunately, this is feasible using reactions such as:

$$n + {}^{6}Li \rightarrow {}^{4}He + T$$
$$n + {}^{7}Li \rightarrow {}^{4}He + T + n$$

Tritium Production

The incident neutrons that induce these reactions arise from the D + T reactions, and perhaps also from other reactions initiated by the primary neutrons. The reaction involving ^{6}Li has the advantage that it is exothermic, adding another 4.8 MeV of kinetic energy to the 17.58 MeV of kinetic energy released by the D + T reaction that produced the neutron. Moreover, because it is exothermic, there is no threshold energy for initiation of the reaction, meaning that even low-energy neutrons which have undergone many elastic and inelastic collisions can still produce tritium in this way.

The reaction involving ^{7}Li is less advantageous in one sense, in that it is endothermic, consuming 2.5 MeV of kinetic energy in order to occur. This also means that only energetic neutrons above about 3 MeV (after allowing for center of mass effects among the reactants) can initiate this reaction. On the other hand, the n + ^{7}Li has the advantage that it does produce an additional neutron, which may initiate an n + ^{6}Li reaction to produce still another triton.

Thus, for a fusion reactor using deuterium and tritium fuel, the raw material for the production of tritium is lithium. The natural abundances of these lithium isotopes on earth are 7.4% for ^{6}Li and 92.6% for ^{7}Li. In order to produce sufficient tritium to at least continuously replenish its fuel supply, a deuterium tritium fusion reactor would be surrounded by a lithium blanket which would produce tritium and capture the kinetic energy of the fusion neutrons and neutron-induced reaction products. As these particles slow down by collisions with the lithium blanket, their kinetic energy will be converted to heat. This heat will, in turn, be used to produce steam to drive electricity-producing turbines. This tritium-breeding blanket may also include other materials such as beryllium. Terrestrial beryllium is 100% ^{9}Be, which can act as a neutron multiplier primarily through the reaction:

$$n + {}^{9}Be \rightarrow {}^{7}Be + n + n$$

This reaction is modestly endothermic, requiring an input of 1.67 MeV in order to occur, meaning that the threshold energy in the laboratory frame for incident neutrons is around 2 MeV.

If, for instance, ^{7}Li is used to make tritium (which is a fairly good approximation to using natural lithium), then with a system efficiency of 33% for net electricity production, 1 kg of ^{7}Li would be sufficient to produce 7×10^{4} GJ of electricity from a D + T fusion power plant. At the present cost of roughly $20 per kg (1) for natural lithium, the lithium in a tritium breeding blanket would contribute only about $0.001 per GJ of electricity. This is negligible compared to the bulk price of electricity of about $20 per GJ. Thus, the price of breeding lithium could rise many-fold before it had a noticeable impact upon electricity costs. In fact, if it appeared advisable to use an isotopically enriched blanket, it would be possible to do so with only a modest effect upon the price of electricity. The energetics are such that the electricity produced with a Li^{6} blanket might be slightly cheaper than with a natural Li blanket, because the extra 4.8 MeV from the n + Li^{6} reaction is much larger than the amount of energy required to enrich the lithium.

The United States has large reserves of easily extractable lithium, mostly in dry salt lakes and saline lakes, which could be produced at prices not greatly different from the $20 per kg of today. This reserve is estimated to be about 5×10^{9} kg (1. With an electricity production efficiency of 33%, this would yield about 3×10^{14} GJ of electricity, an amount that is of the order of 800–1000 times the primary energy consumption of the world in the early 1990s. Inasmuch as the United States comprises only 6% of the earth's land area, it is reasonable to assume that the total world reserves of cheaply extractable lithium might be a few times greater than these United States reserves. Thus, the reserves of easily obtainable lithium are sufficient to run an economy powered by D + T fusion reactors for a period of several to many centuries. There is a much greater amount (of order 10^{3}) of lithium dissolved in the world's oceans than in land deposits (1). This could be extracted at somewhat greater cost than from the saline lakes and dry salt beds. However, this cost premium might be relatively modest. Dikes could be built across tidal flats to isolate seawater which could then be treated the same as saline lakes and salt flats after it had undergone sufficient solar evaporation.

Reduced Activation Fuels

The high-energy neutrons produced by fusion reactors using deuterium or deuterium + tritium as fuels undergo nuclear reactions with the materials forming the structure of the reactor. Some of these reactions result in the production of radioactive nuclei with a variety of half-lives. This activation of structural components, and particularly of those components close to the fusion core, makes maintenance more difficult, and also will require at least short-term storage of removed components, or possibly longer term, depending (as will be discussed in a later section) upon the materials used to build the reactor. In addition, the energetic neutrons can introduce lattice defects as they scatter, which in turn reduce the lifetimes of components, requiring more frequent replacement.

In order to reduce these deleterious effects of the neutrons upon fusion reactor structures, studies have been carried out to evaluate the feasibility of alternate fuels which would produce fewer neutrons per unit of fusion energy released. These reactions include the D + ^{3}He $\rightarrow$ ^{4}He + p reaction and the p + ^{11}B $\rightarrow$ 3 ^{4}He. As discussed earlier, the first of these reactions requires higher temperatures than are presently obtainable, and is more subject to bremsstrahlung radiation losses, and is much more vulnerable to synchrotron radiation losses in systems with magnetic fields than is a D + T fuel mixture. Due to lower fusion cross sections and the higher atomic number of ^{11}B, these problems are more difficult for a reactor using the p + ^{11}B reaction.

A reactor using the D + ^{3}He → ^{4}He + p reaction would have the advantage over a D + T reactor that both components of the fuel would be stable. It would still produce some 3.51 MeV neutrons from D + D reactions, and it would produce some 14.1 MeV neutrons from the burnup of tritium produced in one of the two branches of the D + D reaction. However, almost all of the tritium produced would be consumed in D + T fusion reactions, and thus would not leave the reactor. If the temperature and fueling profiles of such a reactor could be optimized appropriately, the neutron production rate could be reduced by a factor of 100 compared to a D + T reactor (2).

Unlike all the other potential fusion fuels, which occur in abundant commercially attractive concentrations on or near the earth's surface, ^{3}He is extremely rare in the earth's crust. The small supplies that are available arise from the beta decay of tritium. This tritium, in turn, has to be bred through nuclear reactions of one sort or another. At the present time, all substantial production of tritium is carried out through neutron capture reactions on deuterium in heavy water moderated fission reactors. In the future, tritium might be bred through other nuclear reactions either by fusion reactors or high-energy particle accelerators. However, if one must first produce tritium in order to obtain the ^{3}He fuel for a ^{3}He + D fusion reactor, then much of the rationale for using ^{3}He disappears.

What is needed in order that ^{3}He + D fusion reactors have the possibility of commercial viability is a naturally occurring source of ^{3}He in concentrations and quantities sufficient for economic exploitation. As it happens, while no ^{3}He deposits occur on earth, they do occur in abundance on the surface of the moon.

Lunar ^{3}He

Because the moon, unlike the earth, lacks both an atmosphere and a magnetic field, its surface is raked by the solar wind, a flux of energetic particles driven outward from the sun. Among other constituents, this wind carries ^{3}He, which embeds itself in the lunar surface when it strikes. The moon, unlike the earth, does not have enough internal heat to drive plate tectonic motions of its surface now, nor for a very long time into its past. For similar reasons, much of its surface has not experienced fresh outpourings of lava or ash since the interior cooled. The early moon probably possessed an atmosphere, but this long ago escaped into space because of the moon's weak gravity. Consequently, there has been no weather and no water to erode and rearrange the lunar surface in billions of years. When our solar system was young, the moon, in common with the earth, was subjected to an intense bombardment by large and small fragments of matter left over from the formation of the planets. The ejecta arising from this bombardment deposited new layers of material over much of the moon's surface, burying pre-existing surface layers. However, most of the solar system debris was swept out of the planetary orbits long ago through the direct interception of fragments by the planets and moons, and through deflections of the orbits of fragments arising from close encounters with the gravitational fields of the much larger planets and moons. The evidence of this early bombardment has eroded away on earth due to the effects of weather, water, plate tectonic movements of the land, and volcanoes. On the moon, however, which lost these processes in the distant past, many large craters still stand from the later periods of the bombardment.

As a result of these differences between the histories of the lunar and terrestrial surfaces, much of the lunar surface has remained relatively undisturbed for perhaps billions of years. During all this time the solar wind has continued to deposit ^{3}He into the lunar surface, so that it now exists there at a concentration of about 10 parts per billion. A study of the feasibility and economics of collecting and concentrating the ^{3}He from the dust which covers much of the lunar surface found that, even after including the high transportation costs of carrying equipment to the moon and sending the ^{3}He back, this would be economically feasible if a commercially viable ^{3}He + D reactor could be developed (3). As discussed earlier, however, the parameters required for a ^{3}He + D reactor to be viable are significantly more daunting than for a D + T reactor.

Thus, although there are a number of light element isotopes which might someday be suitable as fuels for advanced fusion reactors, the remainder of this article will concentrate on concepts for fusion reactors fueled by deuterium concentrated from water, and by tritium which would be produced by fusion-produced neutrons from lithium in blankets surrounding fusion reactors.

POSSIBLE TYPES OF CONFINEMENT FOR FUSION REACTORS

Because temperatures much higher than those to which we are accustomed are required for any fusion reactor with a useful reaction rate, it is important that there be some method for keeping the reacting fusion fuel out of direct contact with material objects which would quickly cool the fuel to temperatures below which fusion reactions were negligible. The principal mechanism responsible for this rapid cooling would not be conduction into the intruding material, but rather radiation losses from the fuel due to enhanced bremsstrahlung and line radiation from atomic transitions due to heavier impurities entering the fuel. Thus, a reactor requires some sort of restraining force which balances the outward pressure of the fusioning fuel, which, since it is absorbing part of the energy it releases, is also producing pressure. There are a number of restraining forces one might imagine, not all of which are practical in a power plant.

Gravitational Fields

Gravitational fields produce suitably high restraining forces only for very large assemblages of mass, because the gravitational force is much weaker than the other known fundamental forces of nature. Thus, while all the successful fusion reactors we can see, namely stars, use gravitational fields for confinement, this is impractical for a commercial reactor by an enormous factor.

Dc Electric Fields

If the reacting fusion fuel possessed a strong net electrical charge, then it might be possible to confine it with dc electric fields. However, in the types of fuel assemblages so far used in fusion research, any net imbalance in charge which developed in the fuel was much too weak to permit confinement solely through the forces that could be applied through dc electric fields. Significant electric fields do develop under some conditions inside the fuel using the magnetic field confinement approach discussed later, and they may play a role in altering the quality of the confinement when they do occur.

Radio Frequency Electromagnetic Fields

At the high temperatures required for a practical fusion reactor, matter exists not in any of the three states with which we have everyday experience: solid, liquid, or gas, but rather in a fourth state known as a plasma. In this state, the negatively charged electrons have been stripped away from positively charged nuclei which they otherwise normally encircle to form electrically neutral atoms. The plasma consists of unbound electrons and these nuclei, called ions. This has the result that both the ions and the electrons are highly mobile, and can rapidly rearrange themselves to counter external electric fields applied to the plasma. Because electrons are much lighter than ions, they have far higher velocities if they are at a temperature roughly similar to that of the ions. Thus, the higher electron mobility normally accounts for most of the charge rearrangement which takes place in a plasma to shield out externally applied electric fields. Plasma will be discussed in more detail later in this article.

An approach which has been considered for confining a fusioning plasma is to counter the outward pressure of the plasma with the inward radiation pressure of a radiofrequency electromagnetic field. However, for plasmas that would produce fusion power densities appropriate to this type of fusion reactor, the overall outward plasma pressure would be one atmosphere or more. In order to produce a countering inward radiation pressure of about an atmosphere, the electric field strength in the applied wave would need to be about 1 million V/cm. This is difficult to achieve without inducing electrical breakdowns, and eddy current losses due to image charges in nearby walls might be large.

Inertial Confinement

Inertial confinement is the method used in fusion bombs, more commonly called hydrogen bombs. It works well under those conditions when the radiation from the explosion of a fission bomb, usually called an atomic bomb, is used to produce soft X-rays, which isoentropically compress a fuel mixture of deuterium and tritium. The goal of inertial confinement research for fusion reactor applications is to reproduce this effect on a much smaller amount of fuel, and with a far less energetic driver than an atomic bomb. This approach is referred to as inertial confinement because it is simply the inertia of the assemblage which confines the fuel while it is fusioning. The fuel is confined for a time approximately given by the time required for the hot ions, with a mean velocity of about 10^8 cm/s, to traverse the radial dimension of the fuel, which is much less than 1 cm. Thus, the confinement time is in the sub-nanosecond range. The fusion output is proportional to the product of the density, the ion temperature, and the confinement time. This is more conveniently referred to as $n\tau T$, where n is density, τ is confinement time, and T is the ion temperature. For realistic ion temperatures of a few tens of keV, the very short confinement time of the inertial approach requires compression of the fuel to very high densities, 40 to 400 times the normal density of the solid deuterium and tritium fuel (4), in order to produce fusion energy releases relevant to a reactor.

Confinement by Magnetic Fields

An alternate strategy for maximizing the $n\tau T$ product of density, confinement, and ion temperature is to use a much lower density, but a correspondingly longer confinement time. The technique which most naturally fits this approach is to place the plasma in a magnetic field. Since the plasma is composed of electrically charged particles, they are constrained to move in helical paths along the lines of magnetic force, with the negatively charged electrons spiraling in one direction, and the positively charged ions in the other direction. In the simplest instance of a uniform magnetic field, and in the absence of collisions, a charged particle remains tied to its line of force, although it is free to move along it. Consequently, the overall effect of a uniform magnetic field is to restrict the outward motion of particles across magnetic lines of force, while leaving motion parallel to the magnetic field unimpeded. For a plasma with a pressure perpendicular to the magnetic field lines of 100 kPa (1 atm) (an overall pressure that is of the general order required for a fusion reactor), a countering perpendicular magnetic force of 100 kPa (1 atm) can be produced with a field strength of about 5000 Oersteds, which is quite readily achievable.

Principal Confinement Approaches of Fusion Research

Research toward the goal of a fusion reactor began in a significant way in 1951 at Princeton University, and has since spread to many different laboratories in a large number of nations. The overwhelming majority of this research has followed either the low-density, high-confinement time approach using magnetic fields, or the high-density, low-confinement time path of inertial confinement. Accordingly, the remainder of this article will cover only magnetic confinement and inertial confinement, with an emphasis on magnetic confinement because this has profited from the most research, and is presently closer to practicality.

MAGNETIC CONFINEMENT FUSION

Physics of Plasmas

Debye Length. Because a plasma is made up of positively charged ions and negatively charged unbound electrons, these two populations of particles interact strongly with each other through the coulomb force. Any substantial displacement of one species relative to the other leads to a large electrostatic potential, which would require some countering force to maintain. In the absence of any externally applied countering force, there is always an intrinsic force available from random thermal fluctuations in a plasma. The electron thermal energy density per degree of freedom is $0.5n_eT_e$ in a plasma with an electron density of n_e and an electron temperature of T_e. This energy is available to drive charge separations between the positive ions and the negative electrons. The restoring force is provided, in turn, by the electrostatic energy density associated with the electric field established by the charge separation. For a charge separation length of d, the electrostatic energy density E scales approximately as:

$$0.5\epsilon_0 \mathrm{E}^2 \approx 0.5\epsilon_0 (n_e e d/\epsilon_0)^2$$

where ε_0 is the permittivity of free space, and e is the fundamental charge of an electron. If this electrostatic energy density is compared to the thermal energy density, then it is apparent that substantial charge separations within the plasma can only take place over distances not significantly greater than $d \approx \lambda_D$, where the value of λ_D is given by:

$$\lambda_D = (\epsilon_0 T_e/ne^2)^{0.5}$$

This is called the Debye length, because it was first calculated theoretically for electrolytes by Debye and Huckel in 1923. This equation is the one most commonly used for λ_D, although a more precise definition would include a term $(1 + ZT_e/T_i)$ in the denominator to account for the ion effects. Here T_i is the ion temperature and Z is the ionization state of the plasma ions. It is of most importance for plasmas with high charge state ions, which is not the case for the fuel components of a deuterium and tritium fusion reactor, although there is usually some admixture of light impurities such as carbon and perhaps small amounts of oxygen in the plasma of today's experimental fusion devices. This ion term can also be of importance if T_e is much greater than T_i. However, in plasmas of interest to fusion work, T_e rarely exceeds T_i, and is often lower than it.

Since both the ions and electrons in a plasma can move freely, they will tend to move so as to neutralize the electric fields arising from charge imbalances, with electrons positioning themselves to shield the electrostatic field from positive ions, and the ions moving to cancel the field from the negative electrons. The electrons, being much lighter and therefore much more mobile than the ions for all plasmas of interest to fusion, account for most of this charge shielding.

The Debye length is one of the fundamental parameters of any plasma. Charge imbalances can occur over distance of the order of the Debye length, but over distances much greater than the Debye length the plasma will not sustain a net electric field unless there is some additional countering force to support it. Similarly, over distances much greater than a Debye length, a plasma will shield out externally applied electrostatic fields. Thus, the electric potential is normally the same throughout a plasma, unless some other force due, for instance, to rotation, alters the balance.

The two criteria for defining an assemblage of ions and electrons as a plasma both involve the Debye length. The first criterion is that the spatial extent of the plasma should be much greater than a Debye length. The second criterion is that there should be many charged particles within a Debye sphere (with a radius λ_D) so that the statistical treatment underlying the definition of the Debye length is valid. The Debye length is usually small for the plasmas used for laboratory or industrial applications. For a 3 eV arc discharge with an electron density of 10^{19} m^{-3}, the Debye length is about 3×10^{-6} m, with about 10^3 charged particles within a Debye sphere, sufficient to validate the underlying statistical treatment (5).

A plasma more characteristic of those produced in fusion research devices might have a density of 3×10^{19} m^{-3}, T_e = 10 keV, $T_e \approx 0.5T_i$, and $Z \approx 1.2$, with the fact that Z is not unity arising from light contaminants in the deuterium and tritium fuel. For these conditions, the Debye length is about $\lambda_D \approx 8 \times 10^{-5}$ m, where almost all of the increase relative to the arc discharge is due to the much higher temperature of the fusion plasma. Since the characteristic dimensions of fusion plasmas are usually of the order of a meter or meters in magnetic confinement devices, it is always the case that they are much greater than a Debye length.

The plasmas found in electric arcs, lightning, or fluorescent lights are usually only weakly ionized, with perhaps one to a few percent of the gas molecules ionized. On the other hand, plasmas with the much higher temperatures needed for magnetic confinement fusion research are highly ionized, with thermal neutrals only penetrating into the outer few centimeters of the plasma.

Cyclotron Motion in Magnetized Plasmas. If a magnetic field is applied or arises within a plasma, then the motion of the constituent electrons and ions is significantly altered, becoming more ordered along a preferred axis. The equation of motion of a particle with charge Z, vector velocity v and mass m in a magnetic field of vector strength B is given by:

$$m(d\boldsymbol{v}/dt) = eZ(\boldsymbol{v} \times \boldsymbol{B})$$

If we choose the z axis to be along the magnetic field, then, in cartesian coordinates, the components of the particle motion are as follows:

$$dv_x/dt = \omega_c v_y$$
$$dv_y/dt = -\omega_c v_x$$
$$dv_z/dt = 0$$

with $\omega_c = (eZB)/m$, and with B the scalar magnitude of the magnetic field. The fact that the particle velocity is constant along the magnetic field arises from the fact that, as shown in the vector cross product above, magnetic fields have no effect on a velocity component parallel to the field. The x and y components are both perpendicular to the magnetic field. The time derivatives of a particle in these dimensions correspond to circular motion with a frequency of ω_c, which is called the cyclotron frequency. The combination of this circular motion perpendicular to the magnetic field and the uniform velocity parallel to the field (for uniform magnetic fields) give rise to a net helical particle orbit spiraling along the magnetic field, as shown in Fig. 1.

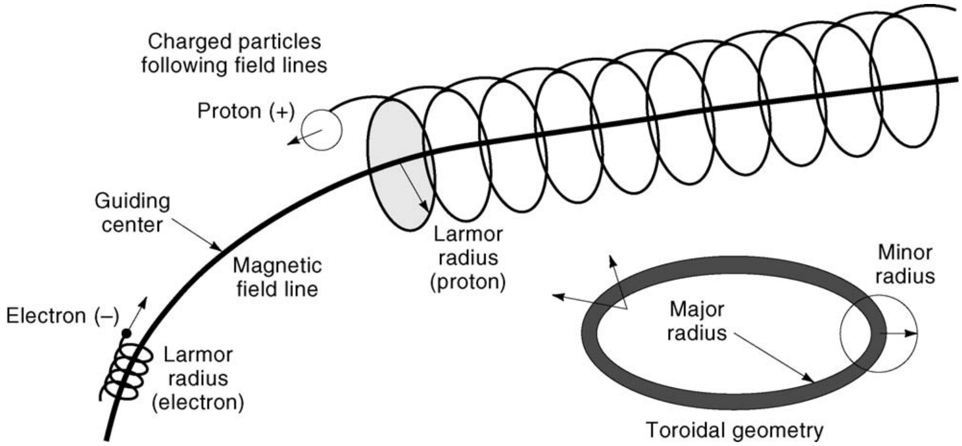

Figure 1. Charged particles gyrating along a magnetic field line. For electrons and protons with the same energy, the Larmor radius for the electrons would be 1/43 of the proton's. The average motion can be described by the guiding center, so long as no parameters change significantly across a Larmor orbit. Drawn by T. Stevenson.

In a plasma with an isotropic Maxwellian velocity distribution, the temperature perpendicular to the magnetic field is equal to the temperature parallel to it. However, because the perpendicular motion encompasses two degrees of freedom while the parallel motion has only one, for these conditions the perpendicular motion of the plasma particle distribution carries twice the energy content of the parallel motion. It is not always the case that the plasma temperature is isotropic with respect to the magnetic field direction. Some heating mechanisms transfer energy preferentially into either the parallel or the perpendicular motion of the particles with respect to the magnetic field direction. Similarly, some particle loss processes involving collisions preferentially deplete either the perpendicular or the parallel energy, depending upon whether the process is stronger for large perpendicular energies or large parallel energies.

Larmor Radius in a Magnetized Plasma. The radius of the perpendicular motion around the field line is called the Larmor radius, denoted as ρ. It is given by the ratio of the perpendicular velocity to the cyclotron frequency, and scales for a particle moving at the thermal velocity as:

$$\rho = [(2)^{0.5} v_t]/[|Ze|B]$$

where v_t is the thermal velocity, and the absolute value of the charge is taken to ensure a positive length. The factor of square root of two arises from the two degrees of freedom in the perpendicular motion. For any plasma, whether Maxwellian or not, there will clearly be a range of values of the Larmor radius corresponding to the perpendicular velocity distribution. Inserting the values of electron and baryon masses leads to the following more specific formulas for the cycloton frequencies and Larmor radii of thermal particles in plasmas: For an electron:

$$\omega_e = 1.76 \times 10^{11} B \,\mathrm{s}^{-1}$$
$$\rho_e = 1.07 \times 10^{-4} T_e^{1/2}/B \,\mathrm{m}$$

For an ion of charge Z and a mass of A atomic mass units:

$$\omega_i = 9.58 \times 10^7 (Z/A) B \,\mathrm{s}^{-1}$$
$$\rho_i = 4.57 \times 10^{-3} (A^{0.5}/Z) T_i^{0.5}/B \,\mathrm{m}$$

In these equations the electron and ion temperatures are expressed in keV. For comparison, room temperature is about 0.025 eV. The magnetic field is given in units of tesla, where a tesla is equal to 10,000 gauss, or roughly 25,000 times the average strength of the earth's surface magnetic field.

In a 5 T magnetic field, which corresponds to a typical field strength on the plasma axis of the Tokamak Fusion Test Reactor at Princeton University, a deuteron would have a cyclotron frequency of $\omega_i = 2.4 \times 10^8$ per second. A plasma electron in the same field would have a cyclotron frequency of $\omega_e = 8.79 \times 10^{11}$ per second. The much higher frequency of the electron cyclotron motion is a consequence of the much smaller mass of the electron relative to a deuteron. For the same reason, the Larmor radius of an ion is much larger than that of an electron. For the case of a 5 T magnetic field confining a plasma with central temperatures of 40 keV for the ions and 12 keV for the electrons, which is reasonably representative of higher performance plasma parameters on Princeton's Tokamak Fusion Test Reactor, the Larmor radius of an electron with the thermal energy would be $\rho_e = 0.0073$ cm, and for an ion with the thermal energy it would be 0.58 cm. The full range of Larmor radii would include values ranging from a bit smaller than this up to a few centimeters (on the tail of the Maxwellian distribution) that would be a factor of 2 or more larger. These Larmor radii are small compared to the dimensions of the fusion devices in which they presently occur. The position of a charged particle averaged over its cyclotron motion is called the guiding center. In the absence of perturbations such as collisions, the guiding center of an electron or ion moves along a line of magnetic force (with the electron guiding centers moving in one direction, and the positive ion guiding centers in the other). So long as the gradient scale lengths for plasma properties such as density, temperature, and magnetic field strength are much larger (usually a factor of a few is sufficient) then it is a good approximation to model most types of particle behavior by following the guiding centers rather than the detailed gyromotion, which is more complicated computationally, and requires more computer time.

Plasma Diamagnetism. The cyclotron orbits of particles in a plasma each enclose small volumes of magnetic flux.

As a consequence of Lenz's Law, the directions of rotation of both the electrons and the ions are such that the tiny solenoidal currents they represent produce magnetic fields in the opposite direction from the field they are enclosing. Thus, the gyromotion at the cyclotron frequency reduces the total strength of the field inside a particle orbit. This is referred to as plasma diamagnetism. Its importance increases as the plasma pressure increases or the externally applied magnetic field decreases. For high ratios of the plasma pressure to the applied magnetic field strength, this diamagnetism can hollow out the magnetic field and produce a region in the central plasma with a magnetic well, that is, a region which is everywhere surrounded by increasing magnetic field strength. This approach has been suggested as a technique for confining plasma, but its practicality is not yet clear.

Magnetic Moment. A particle gyrating around a magnetic field line constitutes an electric current $I = |Zq|\omega_c/2\pi$, which encompasses an area of $A = \pi r^2_L$, where Z is the ionization state, q is the fundamental charge, ω_c is the cyclotron frequency, and r_L is the Larmor radius. The product of the current and the enclosed area is called the magnet moment of the particle orbit, and is generally denoted as μ. Thus,

$$\mu = IA = mv_p^2/2B = W_p/B$$

Here v_p is the velocity of the gyrating particle in the plane perpendicular to the magnetic field line it is following, and W_p is the kinetic energy associated with this perpendicular velocity. It can be shown (5) that, in the absence of collisions or electric fields, the magnetic moment μ of a particle orbit is an invariant. This has far-reaching consequences for plasma confinement schemes utilizing magnetic fields. The invariance of a particle orbit's magnetic moment means that as the particle gyrates along a line of magnetic force into a region of stronger magnetic field, the perpendicular velocity increases so that the energy of rotation perpendicular to the field line increases by the same factor as the magnetic field strength. Inasmuch as the particle's total kinetic energy is also constant in the absence of collisions or electric fields, the increase of the perpendicular energy as a particle gyrates into an increasing magnetic field implies that the kinetic energy of the particle parallel to the field line decreases, and thus that the velocity along the field line also decreases. Thus, for a particle orbit which has a finite energy of perpendicular rotation along any part of its path, there can exist some magnetic field strength at which its velocity parallel to the field line goes to zero. When this happens, the particle is reflected back into the region of weaker magnetic field. This is referred to as *magnetic mirroring*, and plays a significant role in every form of magnetic confinement. If a region of weaker magnetic field is bounded by a stronger magnetic field at each end, then particles can be reflected back and forth between the regions of stronger magnetic field, producing a trap, as shown in Fig. 2. This works for both positively charged particles (most ions) and negatively charged particles (for example, electrons). For a particle with parallel velocity v_z and total kinetic energy W_t, the parallel velocity will vary as:

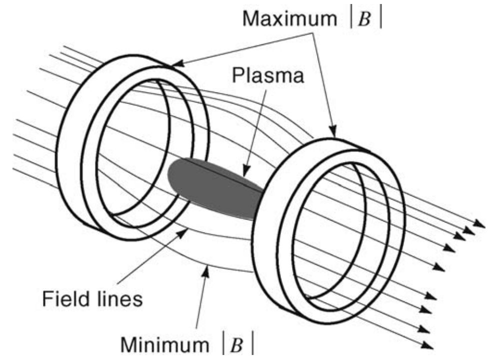

Figure 2. The simplest form of mirror confinement. The trapped portion of the particle velocity distribution reflects from the stronger magnetic fields at each end of the cell, forming a confined plasma in the region between the mirrors.

$$mv_z^2/2 = W_t - \mu B$$

Clearly there is one class of particle orbits for which magnetic mirroring would have no effect at any magnetic field strength, no matter how strong. Particles which are everywhere moving entirely parallel to the local magnetic field have no gyro-orbit, and thus $\mu = 0$. As a consequence, their parallel velocity is unaffected by changes in the magnetic field strength, and these particles are not reflected.

Magnetic Mirror Confinement. More generally, for finite ratios of magnetic field strengths in a plasma confinement device, a considerably broader range of particle orbits is not mirror confined. If we define the minimum magnetic field strength along a line of force to be B_m and the maximum to be B_M, then the constancy of μ and W_t lead to the condition that all particle orbits with $\mu > W_t/B_M$ are trapped by the magnetic mirror field. If this were not so, then particles could reach the point of maximum field strength with their perpendicular kinetic energy greater than their total kinetic energy, which is not possible. Applying this principle to determine the boundary between mirror trapped and untrapped particles, one finds the conditions for marginally trapped particles, where $W_p(\min B)$ is defined as the energy of perpendicular rotation when the particle is in the region of minimum magnetic field strength, and $W_z(\min B)$ is defined as the parallel energy of the particle at the minimum magnetic field:

$$W_p(\min B) = \mu B_m = W_t B_m/B_M$$
$$W_z(\min B)/W_t = (1 - B_m/B_M)$$

These conditions can also be written in terms of the ratios of the perpendicular and parallel velocities evaluated at the minimum magnetic field with respect to the total velocity, v, of the particle, giving:

$$dv_x/dt = \omega_c v_y$$
$$dv_y/dt = -\omega_c v_x$$
$$dv_z/dt = 0$$

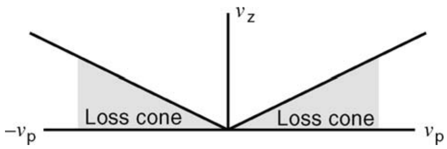

Figure 3. Schematic sample of loss cones in the perpendicular versus parallel velocity space of a magnetic mirror-confined plasma. Particles which undergo collisions such that their v_p/v_z drops into the loss cone quickly escape along the field lines at the magnetic mirror throat.

These equations define the boundary between mirror trapped and untrapped particles in the space of parallel versus perpendicular velocities. Thus, charged particles with a velocity ratio $v_z(\min B)/v$ which is sufficiently low will be trapped. On the other hand, charged particles will promptly escape along the mirror field if they satisfy the criterion:

$$v_z(\min B)/v > (1 - B_m/B_M)^{0.5}$$

This can be rewritten in terms of the ratio of the parallel and perpendicular velocities evaluated at the minimum magnetic field strength:

$$v_z(\min B)/v_p(\min B) > (B_M/B_m - 1)^{0.5}$$

This condition defines two regions in velocity space, in which particles with higher ratios of $v_p(\min B)/v_z(\min B)$ are trapped, and ones with lower values of this ratio are lost. Since the perpendicular velocity corresponds to rotation in two dimensions, the boundary between trapped orbits and loss orbits in a three-dimensional velocity space forms a cone. Thus, particles with orbits that have ratios of parallel velocity to perpendicular velocity which are too great for trapping are said to fall into the *loss cone*, depicted in Fig. 3.

A mirror-trapped plasma can never be isotropic in velocity space. This is because the transit time along a magnetic field line is very fast, so particles with velocity ratios that fall within the loss cone will rapidly escape through the mirror throat (region of maximum B). Thus, the average perpendicular energy in a mirror-confined plasma is larger than the parallel energy by more than the factor of 2 one would expect just from the relative number of degrees of freedom involved.

Note that neither the charge nor the mass of the particles enters directly into the criteria for trapping. In an entirely collisionless plasma, the particles within the loss cone would be entirely lost within a transit time (the time required for the slowest particles to travel the length of the region enclosed by the magnetic mirrors). After this very rapid loss, the remainder of the plasma would remain confined for as long as the magnetic field configuration was maintained.

In all systems of interest to fusion research, this static condition never arises because coulomb collisions between mirror-confined particles alter their ratio of perpendicular to parallel velocity, creating a dynamic loss due to particles which scatter from trapped velocity ratios to untrapped ratios. Once in the loss cone, these newly scattered particles will rapidly escape the confinement device unless their ratio of perpendicular to parallel velocity is scattered back into the trapped region by a subsequent collision or collisions. Accordingly, although the loss cone is the same for different types of particles in a plasma, the particle species which collides most frequently, and thus will most rapidly scatter into the loss cone, is the one which is preferentially lost.

For a plasma in which the temperatures of the electrons and the ions are of the same order, the electrons are the more collisional species, so they initially escape through the mirror throats more rapidly than the ions. This immediately gives rise to a net negative charge leaving the mirror throats, and a net positive charge remaining in the central cell of the mirror confinement device. This charge imbalance produces an electric potential which is positive in the central cell, and thus an electric field which retards the loss of electrons along magnetic field lines through the mirror throats. This electric field causes lower energy electrons to be confined which would otherwise be lost, and, correspondingly, it slightly increases the parallel velocities of positive ions as they approach the mirror throats. Thus, the electric field strength will rapidly build up only to the strength at which it impedes the loss of lower energy electrons sufficiently so that the net electron loss just balances the loss of positive ions. In a mirror-confined plasma it will be the more energetic electrons which are able to escape over the electrostatic potential hill formed at the mirror throats, so the energy balance of mirrors is dominated by these electron thermal losses. The equilibrium rate at which electrons and positive ions escape from the mirror-confined plasma is governed by the rate at which the less-collisional ions scatter into the loss cone, so ion collisions set the rate of particle loss.

During the early decades of fusion research, devices incorporating magnetic mirrors as their primary confinement mechanism were extensively investigated. In their simplest form these machines consisted of a central cell permeated by a uniform solenoidal magnetic field, bounded by a stronger magnetic mirror field at each end. These mirror machines were examples of "open" confinement devices because their magnetic field lines left the region of plasma confinement through the throats of the mirrors. This geometry had a practical appeal because it is mechanically simpler to build a linear device than one with some more complex shape. Moreover, mirrors are devices which do not require a time varying field, and thus are intrinsically capable of continuous operation if the engineering systems are designed accordingly. In a practical mirror fusion reactor the magnetic field coils would be superconducting in order to avoid ohmic dissipation.

Simple mirror confinement machines encountered difficulties due to macroscopic instabilities which rippled the plasma surface and greatly augmented the loss of plasma particles and energy. These problems were addressed by confinement schemes which added additional current-carrying conductors to alter the magnetic field in the central cell so that it increased toward the edge of the plasma. The intrinsically anisotropic velocity distribution residing within a mirror drives another smaller scale length microinstability, called the loss-cone instability, which increases particle and energy loss. Finally, of course, there is the overriding consideration that even if all losses across the magnetic field were eliminated, the loss of plasma particles and energy along the magnetic field lines through the mirror throats would remain.

Even in the absence of macroscopic or microscopic instabilities, the theoretical limit on the energy gain of a simple mirror with a single confinement cell is relatively modest. We can define Q as the ratio of the nuclear power (thermal) produced by the plasma to the amount of power supplied via external means to the plasma in order to heat it to the high temperatures required for fusion reactions. Assuming a conservative efficiency for converting thermal energy to electricity of about 33%, and keeping in mind that a practical reactor must produce considerably more electricity than it uses, an energy multiplication factor of $Q = 15$ to 20 or more is often considered to be sufficient for an economic fusion reactor using a fusion-heated steam cycle to drive the electricity-producing turbines.

Due to the end losses along the magnetic field lines, it would be difficult even in principle to achieve an economically acceptable Q with a simple mirror machine. For a simple mirror, free of any instabilities, and using as fuel deuterium and tritium, the energy multiplication would be at best $Q = 2\log(B_M/B_m)$. The logarithm in this expression can never be greater than an order of unity (that is, with a value of one to at the very most a few). This is because the mirror ratio, the ratio of the maximum to minimum magnetic field in the device, cannot be made arbitrarily large. Such practical considerations as mechanical forces and the maximum current densities superconductors can carry limit the maximum field strength that can be produced in a mirror of useful size, while the minimum field strength cannot be reduced too much without allowing instabilities and diffusion across the magnetic field in the central cell to increase.

In order to circumvent these limitations, researchers contrived more sophisticated mirror confinement schemes such as tandem mirrors. In this approach, the central confinement cell is enclosed by an additional mirror cell at each end which accepts particles escaping through the velocity loss cone of the mirror field of the central cell. The plasma in these end cells, and particularly the electron component, is heated to very high temperatures to increase the magnitude of the electrostatic field plugging the mirror throat from the central cell. Increasing this negative potential in the throat means that more energetic electrons are confined in the central cell, thus reducing the parallel heat loss, and permitting the theoretically obtainable Q to be greater than for a simple mirror.

Maintaining the proper velocity distributions of particles in tandem mirrors without driving instabilities proved quite daunting so that, in the mid-1980s the experiments with the mirror approach to fusion were largely abandoned. Since that time almost all magnetic confinement fusion has gone into the other general class of devices called *closed field systems*, in which the magnetic field lines pervading the central confinement region do not leave the plasma, but instead circle around to eventually reconnect with themselves.

Closed Magnetic Field Line Confinement Devices. Experimental confinement machines with closed magnetic fields have in most cases used a topology which is basically toroidal, or toroids elongated with straight sections. Such devices were pioneered in western countries by Professor Lyman Spitzer at Princeton University in 1951, and at about the same time by researchers at the I. V. Kurchatov Institute in Moscow. Although similar in many respects, the approaches which these laboratories pursued in the 1950s and early 1960s differed in some important ways.

Rotational Transform. All closed field confinement systems are defined by the fact that their magnetic field lines wrap around to reconnect with themselves after some angular translation along the major circumference of the device. The simplest topology for a closed system, and the one which has been most often used for large fusion experiments, is a torus. In this approach, current-carrying coils are arranged around a vacuum vessel so that they produce a toroidal, or donut-shaped, magnetic field along which the confined plasma ions and electrons can gyrate. This geometry is the topological equivalent of deforming a linear magnetic solenoid into a torus by bending it into a circle such that the ends touch. In a toroidal geometry, in contradistinction to the case for a linear system, each of these toroidal field coils has an inner leg and an outer leg with respect to the vacuum vessel and the plasma it contains. It is clear that the inner legs of the toroidal field coils carry the same current as the outer legs, but the outer legs on a torus are much farther apart from each other than are the inner legs, which are often sized so that they abut or nearly abut each other. This gives rise to an asymmetry between the average current density flowing along the inner circumference of the coil system where the coils nestle near each other compared to the outer coil circumference where the coils are widely separated. This assymetry in turn produces a gradient in the strength of the magnetic field produced by the coils; the field is stronger on the inside of the torus, and weaker on the outside.

A toroidal geometry is most naturally described by two dimensions, the major and minor radii of the plasma toroid, where the major radius, R_0, is the distance from the central axis of the coil arrangement to the axis of the confined plasma. The minor radius, r_0, is the distance between the plasma axis and the outside of the plasma cross section. For circular coil systems and plasmas where the ratio of the perpendicular plasma pressure to the inward magnetic pressure is fairly low, the effective major and minor radii are quite similar to what one would get from simply looking at the plasma as a donut. For elongated coils and plasmas, and for high-pressure plasmas which rearrange flux surfaces, these effective radii are somewhat modified.

If the only magnetic field threading a plasma torus is the toroidal field, then the plasma will be only momentarily confined, no matter how strong the toroidal field is made. This arises as a result of the fact that there is a gradient in the toroidal magnetic field strength running from a high level on the inside to a lower level on the outside. This magnetic gradient causes the negatively charged electrons to drift transverse to the gradient in one direction, and the positively charged ions to do the same thing in the other direction. These oppositely directed vertical drifts produce an imbalance in the net neutrality of the plasma above the midplane, compared to the plasma under the midplane. One side will have excess positive charge, while the other will have excess negative charge. These charge accumulations will establish a vertical electric field which is perpendicular to the toroidal magnetic field. In turn, the combination of this electric field and the magnetic field drive a type of particle drift with a drift velocity of $v = (E \times B)/B^2$. For toroids as described here, the direction of this drift is radially outward along the direction of the major radius. This mechanism drives both the negative electrons and the positive ions out together, so no electric field is established which would limit this drift. As the plasma drifts out to larger major radii it quickly runs into the material boundary of the containment device, and is extinguished as impurity influx leads to large line transition radiation losses. Reference 5 contains excellent descriptions of the two types of particle drifts mentioned here, as well as other sorts of drifts which occur in magnetized plasmas.

The fact that a simple toroidal magnetic field cannot confine a plasma is a consequence of the field gradient that necessarily arises from the geometry of any toroid. Thus, there is nothing which can be done to prevent the oppositely directed drifts of the electrons and ions perpendicular to the magnetic gradient.

What can be done instead is to short out the charge imbalance which would otherwise develop by adding a helical twist to the magnetic field lines so that particles gyrating along these lines of magnetic force will spend half of their time above the tokamak midplane and half below the midplane. Since the vertical drift of each species is independent of whether the particle is above or below the midplane, this means that the unidirectional drift of each particle is outward half of the time, and inward the other half, so there is no net change of position.

The helical twist in the confining magnetic field lines is described by the rotational transform, which is the amount a field line moves in poloidal angle (the angle around a cross section of the plasma donut) as it traces itself around the plasma in toroidal angle. The inverse of this quantity, called the safety factor, and denoted as q, is commonly used in describing toroidal magnetic confinement systems.

The local value of q will in general be different for each magnetic flux surface within a toroidal plasma, but it will be uniform on any given flux surface. For flux surfaces with a minor radius significantly smaller than the major radius, which is the case for all of the large devices so far built, the value of the dimensionless quantity q can be well approximated as:

$$q = (rB_\Phi)/(R_0 B_\Theta)$$

Here r is the minor radius of the flux surface being described, B_Φ is the toroidal component of the magnetic field, R_0 is the plasma major radius, and B_Θ is the poloidal component of the magnetic field (the component which imparts the helical twist to field lines on the flux surface). If a field line returns to its starting position after exactly one circuit of the torus, then $q = 1$ on that flux surface. If, for example, 2.7 transits around the torus are required for a field line to return to its initial poloidal location, then $q = 2.7$. For values of $q > 1$ in a circular cross section toroid with the major radius significantly larger than the minor radius, a particle gyrating along a field line experiences a net magnetic well when averaged along its orbit, which tends to abet confinement.

Within a confined plasma of a device of the tokamak type to be discussed shortly, there are arbitrarily many local values of q associated with flux surfaces. However, the class of surfaces with rational values of q often play a special role either in the large-scale internal stability or the gross stability of the plasma. Rational magnetic surfaces are ones for which the safety factor can be represented as $q = m/n$, where m and n are integers, and a field line comes back to its original position after m toroidal and n poloidal rotations around the torus. Rational magnetic surfaces have proven particularly susceptible to magnetohydrodynamic (*MHD*) instabilities, which are large-scale fluid-like perturbations of the plasma and the field lines within it. If they occur on rational flux surfaces well inside the plasma, they can increase the radial loss of energy and particles; if they occur near the outside of the plasma, they can cause the plasma to disrupt.

The magnetic surfaces with low values of rational q tend to be the most pathological, since they are most susceptible to current-driven instabilities. Early researchers found that it was usually easier to maintain gross plasma stability if the surfaces with low rational values of q were buried well inside the plasma, which meant that the q value of the outside was large (as much as 8 to 10 or so). Thus, the dimensionless quantity q came to be called the safety factor, since it was easier to maintain the overall stability of most sorts of toroidal plasmas if the edge q was larger.

Ways to Introduce Rotational Transform. The pioneers of fusion research in 1951 realized from the beginning that a rotational transform would be needed in the field lines of a toroidal or semitoroidal confinement device. Where they differed in the mechanisms they used to produce the rotational transform. The group at Princeton University imposed the rotational transform externally by applying magnetic fields from helical coils spiraling around the vacuum vessel inside the much more powerful toroidal field coils. Alternatively, in some early experiments researchers obtained the rotational transform without the helical coils, but instead with an elongated racetrack type of design distorted so that the two straight sections crossed over each other. These closed field devices in which the rotational transform arose from the external configuration or special coils were named *stellarators* because of the astrophysics background of their inventor, Lyman Spitzer, Jr., of the Princeton University astrophysics department.

Stellarators had two particularly appealing characteristics. One was that they did not require any net current in the plasma to maintain the rotational transform, making them less susceptible to damaging disruptions and obviating the need to find a way to drive the current. Since the required fields were static, a stellarator was intrinsically capable of steady-state operation if the coils were designed for this. The other major desirable feature was that, with the magnetic fields all imposed by external coils, position control was relatively straightforward, so the plasma could be kept out of contact with material surfaces. This meant that keeping the plasma fairly free of impurities, which would dilute the fuel and increase energy losses through radiation, should be feasible.

A disadvantage of stellarators was that the helical coils producing the rotational transform crossed the toroidal field lines produced by the encircling toroidal array of coils. This produced large mechanical stresses which were difficult to accommodate in the 1950s. Partly as a result of this engineering difficulty, no stellarators with large minor radii were built during that period. These early small-cross-section stellarators (typically with minor radii of order 10 cm or even less) achieved disappointing results, which may have been in some part because they were so small and relatively cool (compared to a stellar core) that low-energy neutral particles could penetrate throughout the plasma, undergoing charge exchange with the hotter confined particles which, once neutralized, would escape the confinement. In the late 1960s the stellarator approach was largely abandoned in response to more favorable results obtained at the I. V. Kurchatov Institute in Moscow using the tokamak approach, to be discussed next. In later years, as the understanding of both the physics and the engineering of fusion devices improved, some researchers have returned to the stellarator approach and variations thereon. Better performance was obtained with larger plasma cross-section stellarators built in Germany by the Max Planck Institute, and a much larger stellarator variant is being built outside Nagoya, Japan, by the National Institute of Fusion Studies. These later stellarators have mostly employed magnetic field coils with complex shapes which fulfill the role of both the toroidal field coil and the helical coils of the early stellarators.

Tokamaks. Soviet scientists chose another way to introduce the required rotational transform which employed the toroidal field coils of the stellarator concept, but no helical coils to introduce the poloidal magnetic field component. They instead drove an electric current which flowed toroidally through the plasma. This current produced an encircling poloidal field which vector-added to the toroidal field from the coils to produce a spiraling field with different values of the rotational transform, and thus of q, on each flux surface. The value of the outer surface q, and consequently of the inner surface q's could be easily altered by driving different amounts of plasma current for a given value of the toroidal field. This device was named the *tokamak*, which means something like "magnetic bottle" in Russian. Figure 4 shows the basic schematic of a tokamak.

Inductively Driven Plasma Current in Tokamaks. In all the early tokamaks, and in many later ones, all or nearly all of the plasma current was driven inductively. This was accomplished by adding a transformer solenoid to the tokamak, with the plasma acting as a one-turn secondary winding. As with any transformer, current would continue to be driven in the secondary, that is, the plasma, only so long as the flux produced by the primary winding was changing. Since there is always a practical limitation on the volt-seconds available to drive such a flux swing, and since there also limitations upon the current densities which conductors can carry, inductively driven tokamaks are intrinsically pulsed devices.

In the present generation of tokamaks, these pulse lengths range from a few seconds to a minute. Inductively driven fusion power reactors of the future might have much longer pulses, perhaps of many hours, but they would still have to pause at some regular interval to reset the current in the primary transformer winding. The resulting period during which the fusion plasma was turned off would be short compared to the time when it was making power. Depending upon the thermal inertia of the heat blanket surrounding the tokamak, where heat exchange and tritium breeding take place, the transformer recharge interval might produce either relatively little or significant fluctuation in the electricity output. Perhaps more significantly, pulsed thermal and magnetic field loads are expected to increase the stress on components near the plasma, thereby reducing their lifetimes. As will be discussed later, much progress has been made in finding noninductive ways to drive the plasma current, allowing steady-state operation of future power reactors.

Plasma Limiters in Tokamaks. Another possible disadvantage of tokamaks in their simplest form is that, since the plasma is a current carrying loop, it tries to expand in major radius. This tendency can be countered, however, by adding coils which produce a vertical magnetic field. A more serious concern of early researchers was that simple circular-cross-section tokamak plasmas required contact with a material limiter in order to stabilize their position. This could act as a source of impurities, and would require large amounts of cooling in a power reactor. Most of the more recent tokamak experiments such as DIII-D at General Atomics in San Diego have replaced the limiter with divertors, to be discussed later, which appear to be much more suitable for reactors.

MAJOR COMPONENTS OF A TOKAMAK FUSION REACTOR

Plasma Confinement System

The quality of energy confinement that can be achieved within a tokamak design is the single most important factor in determining the most feasible design for a fusion power plant. If the rate at which energy leaks out of the plasma is too high, then prohibitively large amounts of power will be required to heat and maintain the plasma at the temperatures required to produce useful amounts of fusion energy. The confinement time is characterized by a quantity, τ, which gives the time required for the energy content of the plasma to decline by an e-folding. At the dawn of fusion research, energy confinement times were typically much less than a millisecond. Today they are commonly hundreds of milliseconds, and in some high-confinement plasmas the confinement time is more than a second.

The basic components of the confinement system are the various current-carrying coils which produce those magnetic fields that are externally applied. In addition to the

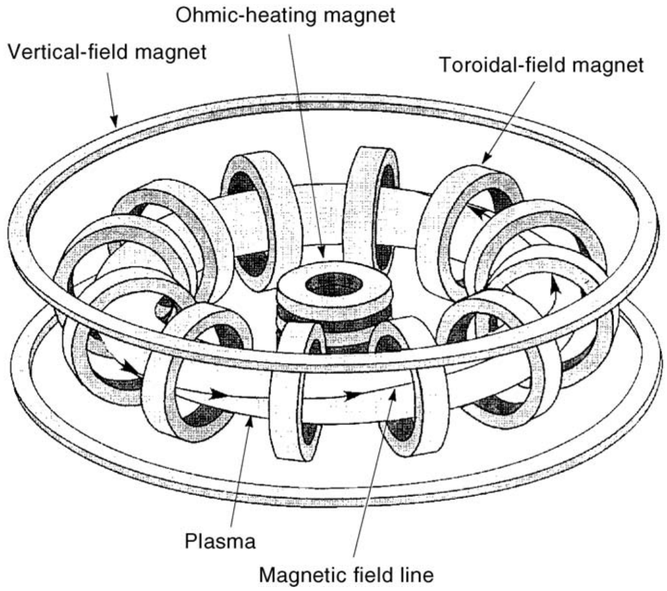

Figure 4. Basic components of a tokamak magnetic confinement configuration. The coils surrounding the plasma produce a toroidal magnetic field, while the current in the plasma creates a weaker poloidal magnetic field. The combination of the two makes field lines which spiral as they are followed around the torus.

toroidal field coils and vertical field coils, there are coils arrayed around the plasma to alter its cross-sectional shape, and in some cases additional coils might be added to control gross instabilities. More recently, the systems which deposit energy, momentum, and additional current into the plasma have also become important in improving confinement.

Collisions among Confined Particles. Getting to the high levels of energy confinement that can be obtained in today's tokamaks has been a long journey. The tokamak confinement configuration is characterized by a magnetic field topology which guides the confined ions and electrons along helical field lines on toroidal flux surfaces. A single ion or electron could travel along such field lines forever if there were no imperfections. In practical applications, there are always some magnetic field imperfections, and, much more importantly, there are many ions and electrons instead of one. The ions and electrons each collide among themselves and with each other. Sometimes when the ions collide they fuse, releasing the energy which is the purpose of the tokamak, but much more often they scatter from each other. The scattering within and among the different confined species gives rise to transport of energy and particles across the confining magnetic field lines, and thus eventually out of the plasma. If this scattering by Coulomb collisions between pairs of particles were the only mechanism by which energy leaked from the plasma, then achieving the quality of confinement necessary for a fusion power plant would have been a daunting, but nonetheless conceptually straightforward, undertaking.

Anomalous Diffusion. The reality has been that the early plasma confinement experiments encountered energy leakage rates much greater than would be expected if the energy loss were primarily due to diffusion arising from two-particle collisions. This type of two-body collisional diffusion was somewhat understood, and could be calculated with neoclassical theory, which took account of effects arising from toroidicity. The excess energy leakage was called anomolous diffusion, and was presumed to be driven primarily by collective instabilities in which great numbers of ions and electrons moved together to form waves and large-scale plasma deformations that transported energy much more rapidly out of the plasma than could two-particle collisions. The early years of fusion research were plagued by anomalous energy losses so large that the confined plasmas could not be successfully heated to and maintained at the very high temperatures necessary for practical fusion power.

Enhanced Confinement. During the decades stretching from the 1970s through the 1990s, progress occurred as successive generations of tokamaks led to a better understanding of how the anomalous energy leakage could be reduced by varying such factors as the dimensions of the plasma toroid and the magnitude of the current flowing within it. At the same time, modes of tokamak operation were discovered with improved confinement: that is to say, significantly reduced anomalous energy leakage. These enhanced confinement regimes were characterized by different profile shapes for the density and temperature than were found in the usual tokamak plasmas, and also by reductions in the rate at which neutral gas was recycling from the walls into the outer part of the plasma. Some of these enhanced modes also involved changes in the current density profile, and in the relative speed at which different portions of the plasma rotated.

The improved modes were usually called by phenomenological names because they were mostly discovered experimentally rather than first being predicted theoretically: fusion research makes progress by building new tokamaks. Two of the most famous confinement modes were the H-Mode, which was discovered in Germany, with the H standing for high confinement to distinguish it from low confinement, and the Supershot Mode, discovered at Princeton in 1986 in the Tokamak Fusion Test Reactor. A supershot plasma was distinguished by having a region in its core with very good energy confinement, and by temperatures much higher than had previously been achieved. Supershot plasmas were also the first ones in a tokamak to exhibit the phenomenon of *bootstrap current.* The bootstrap current arises from a dynamo effect within the plasma as the charged particles press across magnetic field lines. Finding the bootstrap current enhanced the economic viability of future fusion power plants because its existence means that most of the electric current needed to sustain the poloidal component of the magnetic field can be supplied at little or no extra cost by the plasma itself. Similar sorts of plasmas were subsequently produced in Japan's JT-60U tokamak and in Britain at JET, the principal tokamak of the European Union. It was the Supershot mode which was selected as the best route to follow in the Tokamak Fusion Test Reactor to produce fusion thermal power levels of up to 10.7 MW.

Enhanced Reversed Shear Mode. In the mid-1990s, an even better operating mode was found at Princeton's TFTR as well as at the European Union's JET, and subsequently at General Atomics in San Diego, and Japan's JT-60U. This discovery grew out of a study of reversed shear plasmas which had been pursued on several tokamaks. Shear is the rate of change of the rotational transform of the magnetic field lines in passing from one nested flux surface to the next one. It is governed by the shape of the current within the plasma, which produces the poloidal field component that determines the transform. Most tokamak discharges have current profiles that yield shear which is always in the same direction in passing through successive flux surfaces. Techniques were found in recent years which modified the plasma current profile in such a way (by putting more current in the outer portion of the plasma) that the magnetic shear was reversed in the plasma interior. In the central plasma the rotational transform changed oppositely from what it did in the outer plasma in going through successive flux surfaces.

In many cases this reversed shear magnetic field configuration did not produce dramatic improvements in confinement. However, in some cases, the confinement in the central region of the plasma improved markedly in terms of ion energy leakage, and particle confinement of both the ions and electrons. Named the *enhanced reversed shear* mode by researchers at Princeton University's TFTR, who were the first to observe the abrupt transition to this mode with high levels of energetic particle heating power, it was quickly produced on other tokamaks. In retrospect, it was found to be very similar to a type of confinement mode which had been studied earlier in the European Union's JET tokamak with a different heating technique.

Enhanced reversed shear plasmas have nearly perfect particle confinement in their cores, and the rate of energy loss through ion–ion collisions is for the first time down to the theoretical minimum due to two-body collisions. In fact the conduction is so low that it required a revision of the neoclassical theory, which had previously not taken account of orbit effects in regions with very steep gradients in plasma parameters.

At present this enhanced reversed shear regime can only be achieved transiently for periods of roughly 0.2 s to 1 s because the required plasma current profile is produced by rapidly increasing the total plasma current in an already hot plasma. Because the plasma is hot, its electrical conductivity is high, so the extra current takes hundreds of milliseconds or more to diffuse into the central plasma. During this period, the current density in the outer plasma is elevated. It is expected that in the future it will be possible to maintain the required reversed shear profile with active current profile control using injected beams of energetic particles or injected waves. This confinement mode may lead to fusion power plants which are physically smaller and cheaper than had been previously envisioned, but much work remains to be done in understanding how to control the mode, and in how to fuel and remove helium ash from it.

Plasma Heating Methods

Alongside these improvements in energy confinement, equally significant progress occurred in technologies for heating the confined plasma. There are four general ways of heating a magnetically confined plasma. One is with the fusion reactions themselves. When a deuterium nucleus and a tritium nucleus fuse, 80% of the released energy is immediately carried out of the plasma by a neutron which, being electrically neutral, is not confined by the magnetic field that holds the plasma. In an actual fusion power plant, the neutron will be captured in a special blanket where it produces heat for electricity production and new tritium for fuel. The remaining 20% of the energy is borne by the nucleus of a helium atom, called for historical reasons an alpha particle. Since this nucleus does not have any negative electrons bound to it, it carries a net positive electrical charge, and is confined by the magnetic fields in the plasma.

Consequently, the alpha particles produced by fusion reactions remain in the plasma, giving up energy to the plasma particles through collisions until the alphas are thermalized, that is, until they cool down to the temperature of the plasma. The energy transferred from the alpha particles, which are initially much hotter than the plasma, heats it. In an ignited fusion reactor, such as will one day be used to produce electricity, this will be the principal source of plasma heating; in fact the definition of ignition is that the alphas supply sufficient energy to maintain the plasma temperature. However, some other technique will be needed to heat the plasma of even a future electricity-producing reactor to the ignition conditions under which alpha particle heating can take over, and, in any event, the experimental tokamaks in operation today are of too modest a size and capability to reach ignition conditions.

The other three types of plasma heating technologies are called ohmic heating, wave heating, and neutral beam heating. Of these, ohmic heating is the most readily implemented, since it arises automatically from the electric current that flows through the plasma to maintain the poloidal component of the magnetic field which confines the plasma. In much the same way that an electric current flowing through a copper wire heats it as the electrons carrying the current scatter as they move along it, the current moving through the tokamak plasma also heats it.

There is, however, an important difference in how an ordinary electrical conductor and a plasma behave when heated. In an ordinary conductor, such as a copper wire, the electrical resistance rises as the wire's temperature increases, which is to say that the scattering of the electrons increases. Thus, in an ordinary conductor, if the current passing through it is kept constant by raising the driving voltage as the conductor's temperature increases, then the current will become progressively more effective in heating it until the wire eventually melts. In the sorts of plasmas of relevance to fusion, the behavior is just the opposite. The frequency of scattering decreases as the temperature increases, so the electrical resistance drops. This has the consequence that, as the plasma gets hotter, more and more current must be run through it to achieve smaller and smaller additional increases in the temperature. It would be extremely difficult to heat a plasma to thermonuclear ignition temperatures in this way, and even if one could handle the large current required, the poloidal field it would produce might be greater than the optimum for maintaining plasma stability and confinement. In light of these limitations, ohmic heating serves only as the initial heating mechanism in tokamaks, typically raising the temperature to 10 million to 20 million °C, modest by fusion standards, during the startup phase of a tokamak pulse.

One or both of the other two heating techniques—waves or neutral beams—must be used to further raise the temperature to the point where significant numbers of fusion reactions can occur. Wave heating works in much the same way that a microwave oven does, except that instead of heating food by causing molecules to vibrate, the waves increase the energy of the ions or electrons gyrating along the magnetic field within the tokamak plasma. Various forms of wave heating have been used in many tokamak experiments, and it is expected that wave heating will be important for future reactors. However, so far the highest temperature and fusion power results in tokamaks have been brought about with the other technology: neutral beams.

Neutral Beam Injection. Most neutral beams in use today are born in ion sources where the negatively charged electrons are stripped from the positively charged atomic nuclei to produce ions. These ions are then accelerated by passing them between grids with different electric voltages applied to them. This forms a beam of energetic ions moving toward the tokamak, in much the same way that the electron source in the back of a television picture tube forms a beam of energetic electrons moving toward the phosphor screen to form the image. The ions in a beam at Princeton University's TFTR tokamak, however, have more energy than the electrons in the beam of a picture tube—120,000 V instead of a few tens of thousands of volts. In addition, the electric current in the ion beam is thousands of times greater than the current in a picture tube.

The electrically charged ions that emerge from the accelerator grids would not be able to enter the tokamak plasma in their charged state. This is because the very magnetic field which confines the plasma would bend the ion beam out of its path and prevent it from entering the plasma. To circumvent this difficulty, the ion beam is passed through a neutralizer cell filled with low pressure gas, where a portion of the ions each pick up an electron from the gas molecules. After picking up an electron, a beam ion becomes electrically neutral, and thus once again an ordinary atom. Unlike an ordinary atom, which at room temperature has an energy of about one fortieth of a volt, these atoms in what is now a neutral beam have energies of 120,000 V. The remaining ions in the beams are bent out of it with a magnet, and then the purely neutral beam is able to pass unimpeded across the tokamak's magnetic fields to enter the plasma. Once inside the plasma, the neutral atoms of the beam are again ionized through collisions with the plasma particles which detach the electrons from the beam atoms. The resulting 120,000 V ions, being electrically charged, are confined by the magnetic field of the tokamak. They circulate along the magnetic field, colliding with the plasma particles. Since the beam ions are much more energetic than the plasma electrons and ions, they transfer energy to them, and thereby heat the plasma. Eventually the beam ions slow down and become part of the bulk plasma. At this point they are said to have been *thermalized* because their energy is similar to that of the bulk, or thermal ions.

In 1973, when neutral beams were selected as the principal heating technique for the Tokamak Fusion Test Reactor, this technology had reached the point of development where it could inject a few tens of thousands of watts of power into a tokamak plasma. In the experiments which took place at Princeton University in the 1990s, the neutral beams injected a maximum of forty million watts of power into the Tokamak Fusion Test Reactor, an increase of about a thousand-fold in the power capability of the technology from the time when the decision was made to use it.

The fact that the beam ions do become part of the confined thermal plasma means that it is important not only that they carry in energy to heat the plasma, but also that they be the right hydrogen isotopes to fuel it as well. Thus, the TFTR plasma heating systems inject high-energy neutral beams of both tritium and deuterium in order to maintain the correct fuel mix in the reacting core of the plasma.

As the plasmas in successive generations of tokamaks become larger and denser, there is a corresponding increase in the beam energy required to ensure that most of the energy and fuel are deposited in the central plasma. However, the efficiency with which positive ions can be converted back to neutral atoms is a strong function of the ion velocity. For velocities corresponding to a beam energy above 120 keV for deuterium or 180 keV for tritium, the neutralization efficiency is steeply declining into unacceptably low values.

In response to this limitation, a new technology is being developed based upon ion sources which produce negative ions, or ions which have one more electron than the neutral atom would have. Production of negative ions of deuterium or tritium is much more difficult than is producing their positive ions, but the neutralization efficiency of high-energy negative ions is nearly independent of energy across a range of several million electron volts, with neutralization efficiencies of 58% to 59% being quite feasible. The first generation of a negative ion beam system began operating on the JT-60U tokamak at Naka, Japan, in 1996. When it reaches full power it will inject 10 megawatts of 500 keV neutrals. Another large negative ion system will go into operation in the late 1990s on the Large Helical Device, a type of stellarator, outside Nagoya.

Current Drive. All of the tokamaks which have been operated through the mid-1990s normally drove most or all of their plasma current inductively. However, as discussed earlier, this technique necessarily requires that the tokamak plasmas operate in pulses. In order to run future tokamaks steady state or in very long pulses, other techniques must be used to drive the current. Over the years, a number of techniques have been demonstrated to drive substantial amounts of plasma current, or even all of it in some special cases. The methods have used either the high-energy neutral beams which also heat the plasma, or waves which either transfer momentum preferentially to one of the confined species in one direction along the magnetic field lines or, alternatively, waves which preferentially heat the electrons or ions in the direction perpendicular to the magnetic field, thus changing their collisionality. The different current drive methods tend to drive current in different regions of the plasma, meaning that appropriate combinations of them can modify the shape of the current profile to achieve better stability and less energy leakage.

Bootstrap Current. All of the continuous direct current drive techniques described above have in common that their efficiency for driving current is significantly less than the efficiency of neutral beams or some sorts of waves in heating the plasma. Thus, it would be more appealing if much of the plasma current could be driven by the thermal energy of the plasma, rather than being driven by a less efficient externally applied drive. A thermal energy drive would be fed primarily by the efficient neutral beam or wave systems heating a sub-ignited plasma, or by the energy of the fusion-produced alpha particles in an ignited plasma.

A current drive mechanism arising from the bulk thermal plasma was theoretically predicted in the 1970s, and was first identified in the early 1980s on the octupole (a type of magnetic confinement device) at the University of Wisconsin. It was named the *bootstrap current* because it was generated by the plasma itself through a sort of dynamo effect as particles press against magnetic field lines. The strength of this effect depended upon the plasma pressure gradient. Thus, it was not until 1986 that bootstrap current was found experimentally in a tokamak when high confinement plasmas with steep gradients were produced on the TFTR tokamak at Princeton. The existence of the bootstrap current was subsequently verified on the other major tokamaks of the world, and it is now expected to supply much of the current for future tokamaks. One or more of the other direct-current drive techniques may also be needed to provide lesser amounts of localized current drive to produce the optimum profile for confinement and stability.

Because the strength of the total magnetic field in a tokamak increases sharply as one progresses inwards in major radius, particles with sufficiently small ratios of parallel velocity to perpendicular velocity with respect to the field will become trapped in local magnetic mirrors, much as happens to the main plasma in mirror confinement machines. These particles are referred to as *trapped particles*, in distinction from the *passing particles* which circulate freely along field lines around the tokamak. Their orbits, when projected onto a plasma cross section, resemble bananas, with the tips corresponding to mirror reflection from higher magnetic field strengths. In the Soviet Union, these were sometimes referred to as sickle orbits. In lower temperature collisional plasmas, the trapped particles often scatter out of trapped velocity space before they complete a mirror bounce, and thus have relatively little effect on the plasma behavior. At the high temperatures typical of today's tokamaks, the particle collisionality is much reduced, and trapped particles may complete many bounces. It is these trapped particles which are particularly important in the production of bootstrap current. Trapped particles in a tokamak are an effect which arises from the toroidicity, since this is what creates the higher magnetic field at smaller major radius. Thus, the fraction of the total particles which are trapped is proportional to r/R, ratio of the minor radius to the major radius, which is the inverse of the aspect ratio. Thus, in order to maximize bootstrap current, which is in some sense free, reactor designs are sometimes driven toward lower aspect ratios, in the general vicinity of about 3 or less. However, this is somewhat offset by the fact that the fraction of the current that can be driven by bootstrap effects is in part determined by $(R/r)^{0.5}$. In addition, there are also other constraints, such as the scaling of confinement in some plasma regimes, which may drive the optimum aspect ratio to larger values (6, 7).

Comprehensive design studies of a number of variants of tokamak reactors for power plants have been carried out at the University of California at Los Angeles, resulting in a number of designs dubbed with the name *Aries* (8). These for the most part anticipate providing much of the plasma current with the bootstrap effect, and the balance with one or more of the direct-current drive techniques.

Plasma Exhaust System. A long pulse or steady-state tokamak reactor requires something better than a simple limiter to handle the power coming out the edge of the plasma, and to dispose of the helium ash which results from fusion reactions. This is done by altering the closed field line topology of the normal tokamak configuration. By placing an additional shaping coil in the vicinity of the plasma top, bottom, or in both locations, it is possible to redirect the outer field lines so that a thin layer of the plasma on the edge of the plasma cross section is diverted out of the main plasma chamber into a divertor chamber where the plasma, and the helium ash it entrains, are converted to neutrals through a combination of encountering a region of higher neutral particle density and of directly striking plates of carbon or some other material. The divertor region is baffled from the main chamber to reduce reflux, and much of the neutralized plasma outflow can be pumped away as gas. Figure 5 displays one possible divertor configuration.

Ensuring that the divertor plates have an economically attractive lifetime is a significant engineering problem which is not yet fully solved for fusion power reactors; the power density in the scrapeoff plasma flowing toward the divertor will be high, causing thermal problems, and

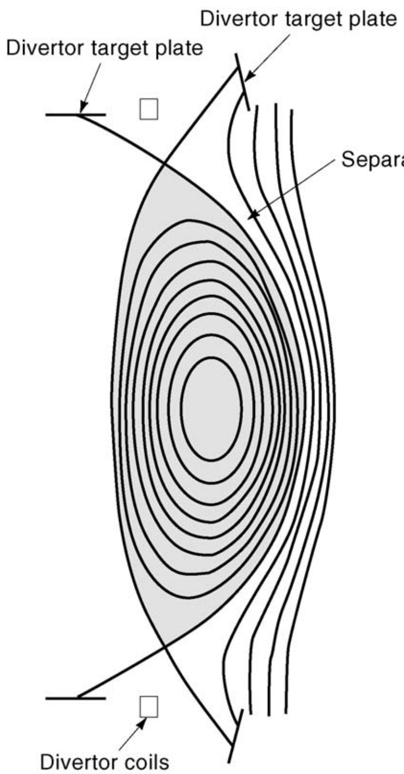

Figure 5. Example of tokamak plasma with top and bottom divertors. The divertor coils produce open flux surfaces after the separatrix field line, allowing the outer plasm to leave the main chamber to strike target plates. This reduces the influx of impurities to the main plasma, allows the power density striking the target plates to be reduced through the expansion of flux tubes, and may allow further power density reductions through radiation in the divertor region.

the energetic particle flux can cause high sputtering rates on the divertor plates. Some of the approaches being pursued to ameliorate these problems include reducing the incident power density by reducing the angle of incidence, and by expanding the flux envelope of the diverted plasma, and reducing the sputtering rate by introducing strongly radiating heavier impurties such as noble gases into the divertor chamber plasma to lower the temperature of the ions so they will produce less sputtering. This deliberately enhanced radiation can also be used to disperse energy at reduced intensity over a larger expanse of material.

Tritium Breeding and Thermal Conversion Blanket. About 80% of the energy produced in a deuterium-tritium fusion reactor immediately leaves the plasma in the form of 14.1 MeV neutrons, since plasmas with dimensions and densities appropriate to reactors are almost perfectly transparent to high-energy neutrons. The remaining 20% of the energy carried by the fusion alpha particles heats the plasma, and eventually either leaks from the plasma as electromagnetic radiation that will be absorbed and converted to heat by the first wall, or it flows in the scrapeoff plasma to the divertor, where it heats surfaces. Thus, all of the alpha particle energy is ultimately collected on surfaces. These surfaces will be cooled with some fluid or gas, of which lithium or helium are examples of possible choices. This extracted heat will then be used to produce electricity, probably by generating steam to drive a conventional turbine.

The neutrons, which carry the other 80% of the fusion power, will be slowed down and captured in a blanket surrounding the tokamak. In the process, their kinetic energy will be converted to heat, which will be further augmented by exothermic nuclear reactions with ^{6}Li. This heat can then be removed from the blanket to produce electricity, either through steam turbines, or possibly by more efficient techniques taking advantage of the fact that the blanket may contain a liquid metal.

This blanket must also breed enough tritium to at least replace the amount consumed in nuclear reactions by the plasma, and it is preferable that it produce somewhat more in order to start up additional fusion power plants. The reactions of neutrons with lithium isotopes which produce tritium were discussed in an earlier section. Proposed blanket designs have of necessity included either natural lithium or ^{6}Li, but they have differed in what neutron multipliers were added to the mix. The most commonly assumed multiplier is beryllium, which undergoes an $(n,2n)$ reaction, but heavy multipliers such as lead have also been considered. The chief difficulty with some heavy multipliers is that they tend to produce radioactive daughter nuclei with undesirable half-lives. Beryllium does not pose this difficulty.

Maintenance and Materials. Developing techniques to maintain and repair fusion power plants is an area which still requires substantial engineering development. It may be somewhat simplified by the time commercial plants begin operating some decades from now by advances in robotics. Present tokamaks are commonly made of stainless steel alloys or inconel, primarily because they are relatively inexpensive, and can be machined and formed by drawing upon a large body of fabrication experience. These materials, however, would stay significantly radioactive for many years after removal from a fusion plant, meaning they would require storage. They would not be volatile or very reactive, so there would be little threat of contaminating water or air.

However, one of the advantages of fusion relative to fission is that one has some degree of freedom to choose materials which reduce activation. This is not possible with fission, since the nuclear waste products are produced by the fissioning of the fuel itself. Accordingly, fusion power plants probably will use different materials which will have lower initial levels of radioactivity per gigawatt year of exposure, and shorter half-lives which decrease the storage time required for removed components. Two of the low activation materials which have been considered by studies of the Aries group (8) include a vanadium alloy with small amounts of chromium and titanium, and silicon carbide. In their reactor models, they found that the total radioactivity per watt of reactor thermal power remaining in a fusion power plant one year after shutdown would be a factor of 10^3 lower than for a comparably rated fission breeder reactor, if the fusion plant was built of HT-9, a ferritic steel. If the plant were instead constructed of the vanadium alloy, the one-year radioactivity was depressed to a level about

10^5 less than for a fission breeder, and if silicon carbide was used for the fusion reactor, then the one-year radioactivity of the fusion reactor declined further to a value about 10^8 less than the fission breeder. If D-^{3}He reactors ever proved feasible through lunar mining, then the silicon carbide design would give a one-year radioactivity that would be more than 10^{10} less than for a fission breeder.

These and other low-activation materials require engineering validation through tests exposing them to large neutron fluences. In addition, some materials, especially silicon carbide, require progress in fabrication technology. The structures normally made of silicon carbide are significantly smaller than would be required for a reactor.

INERTIAL CONFINEMENT

Parameter Regime

A practical fusion power plant fusing deuterium–tritium fuel requires temperatures of over 10 keV, with a fuel density, n, and an energy confinement time, τ, such that their product lies in the range of $n\tau = 10^{14}$ to 10^{15} s/cm^3. A practical fusion reactor using the magnetic confinement approach discussed earlier will likely operate in a regime with a density of about 10^{14} fuel nuclei/cm^3, and a confinement time of about a second or so. Inertial confinement operates at the other extreme of the parameter range. Because the small nuclear explosions in an inertial confinement reactor would occur with high-velocity fuel particles traversing very small distances, the confinement time would be of the order of 10^{-10} s, requiring that the target fuel be compressed to about 10^{25} fuel particles/cm^3.

This compression must be carried out by a strong implosion. A typical inertial confinement target capsule is a small sphere formed of an ablator material which is lined with solid deuterium–tritium fuel. The central spherical cavity thus formed, which composes most of the target capsule's volume, is filled with deuterium–tritium gas. The implosion occurs when the ablation shell is illuminated with a brief burst of extremely high-energy density, on the order of 10 TW/cm^2. This causes the outer portions of the ablation shell to leave it with high momentum which, on the average, will be directed radially outward from the shell. Conservation of momentum requires that an equal inward-directed momentum be imparted to the remaining target, driving the collapse. For a reactor to have energy gains of economic interest, it is also necessary that this implosion proceed isoentropically, with as little preheating ahead of the compression wave as possible. This, in turn, requires that the compression be uniform to a level of about 1% over the entire sphere. Figure 6 shows one possible design for a target capsule.

Direct and Indirect Drive

The most obvious way to compress a target sphere is to shine the driver energy directly onto the ablation sphere. The driver might be an array of lasers, light ion accelerators, or heavy ion accelerators. One of the principal challenges in implementing this approach is that obtaining the required high degree of uniformity in the driver power at the sphere has proven very difficult. In addition, with some types of drivers there may be preheating of the fuel due to penetration by driver ions, and laser-driven instabilities may degrade the symmetry of the compression. Ways of ameliorating these problems are being studied.

The other implosion method is indirect drive. In this approach the fuel capsule is not struck by the primary driver beams (4). Instead, the capsule is placed inside a Hohlraum which is of considerably larger dimensions than the capsule, as depicted in Fig. 7. This Hohlraum is made of a high atomic number material, and has openings through which the driver beams can enter so as to directly strike the inner wall of the Hohlraum, but not the fuel capsule. A portion of the driver energy striking the Hohlraum is converted to soft X rays, and they in turn drive the implosion of the fuel capsule. An advantage of this method compared to direct drive is that quite uniform irradiation of the capsule with the soft X rays can be achieved even with anisotropic primary driver beams, enabling fewer beams to be used. In addition, there is a large body of experience available with this sort of indirect drive from the nuclear weapons program. The majority of research being conducted on inertial confinement in the US today follows the indirect drive approach.

Inertial Confinement Drivers

Laser Drivers. Over recent decades, a number of classes of drivers have been considered, and in some cases tested, for inertial confinement applications. Experiments have been driven by neodymium-glass lasers, and krypton fluoride gas lasers have been studied as possible primary drives. Most of the laser drivers suffer from low efficiencies for conversion of electricity into laser light. In the case of KrF this efficiency is about 6% to 8% (9). This would require that the target capsule produce 140 to 160 times more energy than was in the driver in order to make a power plant feasible. Such high gains are vastly beyond current achievements, but may, in principle, be achievable. Laser diodes have much higher efficiencies of as much as 60%. If the prices of these can be reduced by a large factor, and if techniques can be found to interface them with the thermonuclear environment of a reactor cell, then these may become attractive.

Ion Drivers. The other classes of drivers are light ions and heavy ions. In the light ion approach, extremely high currents of lithium ions are produced by an array of diode sources and accelerated to energies of hundreds of keV to a few MeV. Lithium diode drivers have been studied at the Sandia National Laboratory. The other method would use lower currents of ions from the high mass end of the periodic table such as xenon or cesium, which would be accelerated to energies of 2 GeV to 10 GeV. About 10^5 A of beam at these energies would be required to achieve the necessary illumination intensity. Although the required pulse duration is very brief, such currents are well beyond the capabilities of existing high-energy accelerators. High-current accelerator concepts such as inductive linacs are being studied.

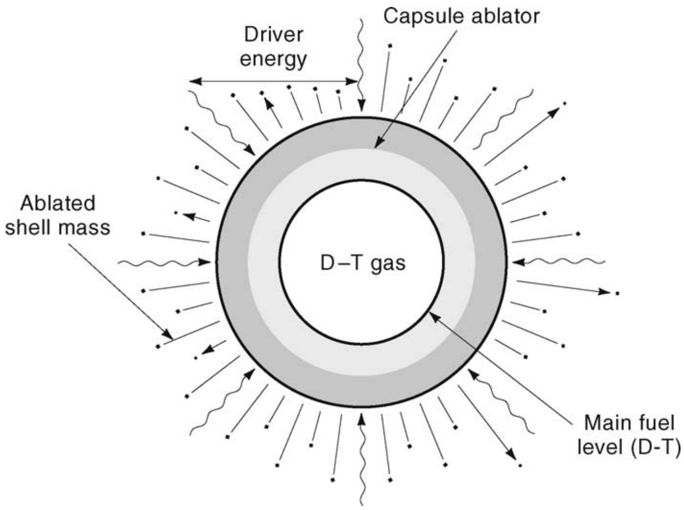

Figure 6. Target capsules for inertial confinement fusion are compressed by ablation from the outer shell. After ignition in the center, the fusion burn region will need to spread outward to produce significant energy gain.

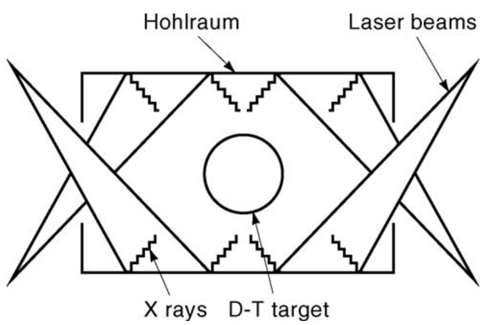

Figure 7. Target placement for indirect drive inertial confinement fusion. High-power laser beams strike the inner surfaces of the Hohlraum producing X rays which compress the target capsule. A variant of this would use heavy ion beams striking absorbers at the ends of the Hohlraum. Indirect drive requires much less perfect symmetry in the drivers than does direct drive.

Both of the ion drive techniques have the major advantage, relative to most of the laser approaches, that their efficiency for conversion of electricity to ion energy can be quite high, 25% to 30%. With the heavy ion approach a target energy gain of a factor of 40 might be sufficient for a feasible plant, which is roughly a factor of 4 below that required for a KrF laser driver. On the other hand, for either technology the final path of approach of the ion beams toward the Hohlraum will be after the final focusing elements of the ion optical system. Maintaining the necessary illumination intensity and uniformity will be difficult as these high space charge density beams converge toward even higher concentrations of space charge. Space charge effects, which arise because of the coulomb repulsion between like charges, are strongly defocussing. In recent years, the vapor pressure in the most common target chamber concepts is much higher than in earlier scenarios. This may result in sufficient plasma being produced along the beam trajectory to partially or totally cancel the effect of the beam space charge forces, and thus simplify the problem of obtaining small beam diameters at the target.

Inertial Confinement Plant Systems

An inertial confinement fusion power plant would consist of a target chamber, the driver system of ion accelerators or lasers, and a target fabrication system introducing several capsules (and Hohlraums if needed) per second into the chamber. A lithium layer would breed tritium and would convert the fusion energy to heat, which could then be converted to steam to drive turbines. Concepts for this lithium layer have included liquid metal flowing over the inner surface of the target chamber, as well as jets of liquid metal around the target.

PROGRESS IN FUSION ENERGY

Over the past four decades, large advances have been made in understanding the physics of plasmas suitable for a magnetically confined reactor, and the technologies for heating the plasmas have similarly progressed. In this period the power released through fusion reactions has increased a factor of 10^8, with over 10 MW produced by the TFTR tokamak funded by the US Department of Energy at Princeton University. Similarly the energy leakage of the ions has dropped from being many times higher than that theoretically predicted to values which in some cases are at or near the theoretically best confinement that could ever occur. Temperatures in tokamaks have climbed from a few hundred electron volts to as much as 40 keV to 45 keV in TFTR and in JT-60U in Naka, Japan. Energy confinement times have grown from a few milliseconds to a second or more on Japan's JT-60U and Europe's JET tokamaks. The power of neutral beam systems heating the plasma has climbed from the tens of kilowatts prevalent a four decades ago to 40 MW on TFTR and JT-60U. In recent years, neutral beam heating systems based upon negative hydrogen and deuterium ions have been deployed in Japan on the JT-60U tokamak in Naka and the Large Helical Device at Tokai. Negative ion based beam systems can maintain good neutralization efficiency at the much higher megavolt beam energies which will be needed for magnetically confined fusion reactors. The JT-60U negative ion based neutral beam has operated at over 400 kilovolts.

New superconducting tokamaks are nearing completion in China, India, and South Korea, while a compact stellarator is being built at Princeton in the U.S. to study a configuration combining some aspects of both stellarators and tokamaks. A coalition consisting of the European Union, Japan, the United States, Russia, South Korea, China, and India is building the International Tokamak Experimental Reactor (ITER) in Caderache, in the south of France. This will be the first fusion device to demonstrate useful amounts of fusion power, and will test many of the technologies necessary for a real fusion reactor. It is expected to begin operations in the latter part of the second decade of this century.

Inertial confinement fusion has made advances in understanding the physics of matter at the very high pressures and densities required for this approach, and understanding has been gained about the ways in which energy and momentum are transferred to the target capsule. A major new laser-driven test bed, named the National Ignition Facility, is nearing completion at Lawrence Livermore National Laboratory in California. It is expected to produce target ignition early in the second decade of this century. Meanwhile, investigations continue on more efficient and compact heavy ion driver beams for inertial confinement fusion reactors.

Producing an economically attractive fusion power plant will require more work for either the magnetic or inertial confinement approaches. Based upon the present state of these fields, it appears that the magnetic confinement approach enjoys a clearer path forward to a reactor, but this could change in the future depending upon progress in inertial confinement.

BIBLIOGRAPHY

1. J. Wesson *Tokamaks*, Oxford: Clarendon Press, 1987.
2. J. G. Cordey R. J. Goldston R. R. Parker Progress toward a Tokamak fusion reactor, *Phys. Today*, 22–30, 1992.
3. L. J. Wittenberg J. F. Santarius G. L. Kulcinski Lunar source of ^{3}He for commercial fusion power, *Fusion Technol.*, **10**: 167, 1986.
4. J. Lindl Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain, *Phys. Plasmas*, **2** (11): 3933–4024, 1995.
5. R. J. Goldston P. H. Rutherford *Introduction to Plasma Physics*, Bristol and Philadelpia: Institute of Physics, 1995.
6. L. R. Grisham *et al.* The scaling of confinement with major radius in TFTR, *Phys. Rev. Lett.*, **67**: 66–69, 1991.
7. L. R. Grisham *et al.* Scaling of ohmic energy confinement with major radius in the Tokamak fusion test reactor, *Phys. Plasmas*, **1**: 3996–4001, 1994.
8. R. W. Conn *et al.* Economic, safety, and environmental aspects of fusion reactors, *Nucl. Fusion*, **30**: 1919–1934, 1990.
9. W. J. Hogan R. Bangerter G. L. Kulcinski Energy from inertial fusion, *Phys. Today*, September, pp. 42–50, 1992.

LARRY R. GRISHAM
Princeton University, Princeton, NJ