Hot fusion

Despite more than 50 years of effort, today's nuclear-fusion reactors still require more power to run than they can produce. **Steve Cowley** says the next step is to get the fusion plasma to generate its own heat – to make itself hotter than the centre of the Sun

It has to be one of the greatest public lectures in the history of science. Indeed, the presidential address by Arthur Stanley Eddington to the 1920 meeting of the British Association in Cardiff is still worth reading for the simplicity and clarity of the arguments alone. But it is his extraordinary vision that stands out nearly a century later. Until Eddington's lecture, it was widely accepted that the Sun was powered by gravitational contraction, converting gravitational potential energy into radiation. Some 60 years earlier, Lord Kelvin had argued that this mechanism means that the Sun can be no more than 20–30 million years old. But using simple arguments based on a wide range of observations, Eddington showed that the Sun must be much older than Kelvin's estimate and that stars must draw on some other source of energy.

It was fortunate that just prior to Eddington's address his Cambridge University colleague Francis Aston had measured the masses of hydrogen and helium to be 1.008 and 4, respectively. Eddington argued that the Sun is being powered by converting hydrogen to helium – by combining four hydrogen nuclei (protons) with two electrons and releasing energy in the process. The exact details were wrong of course – the process is more complicated and involves deuterium, positrons and neutrinos, for example – but the basic idea was correct: the Sun is indeed converting hydrogen to helium.

The energy released in this transformation can be calculated using $E = mc^2$ and the measured masses of hydrogen and helium. From this, Eddington estimated that the Sun has enough energy to shine for 15 billion years – remarkably close to modern estimates of approximately 10 billion years from formation until the Sun enters its red-giant phase, when it will have exhausted the hydrogen fuel in its core. He had deduced the existence of what we now call nuclear fusion. Although Eddington cautioned about being too certain of his conclusions, he realized that the potential was staggering and he immediately saw the enormous benefits fusion could bring society. As he told his audience in Cardiff, "we sometimes dream that man will one day learn how to release it and use it for his service".

Eddington's vision is now within our reach, although it has not been easy getting this far. Along the way we have needed to develop the field of plasma physics, which studies gases heated to the point where the electrons separate from their atoms. Despite the struggles, it is fair to say that scientists have now captured the Sun's power.

Steve Cowley is Chief Executive Officer of the United Kingdom Atomic Energy Authority at the Culham Centre for Fusion Energy and professor of plasma physics at Imperial College London, e-mail steve.cowley@ccfe.ac.uk

1 Deuterium–tritium fusion

deuterium

helium (3.5 MeV)

neutron (14.1 MeV)

tritium

helium

lithium-6
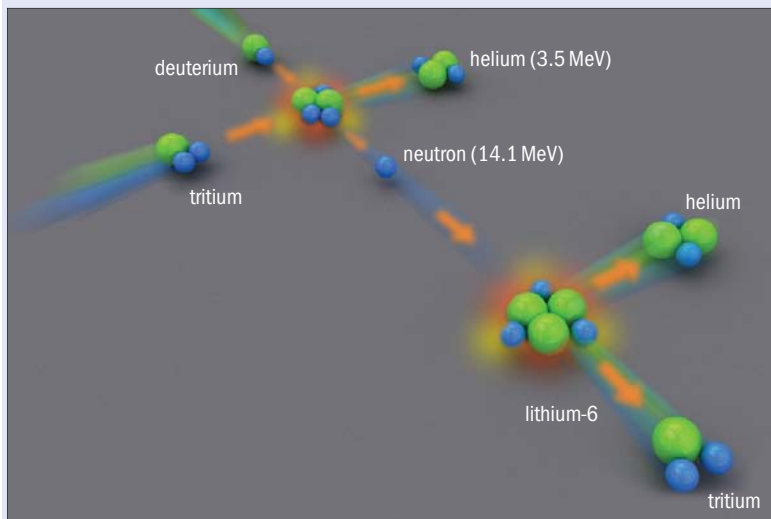
tritium

Deuterium (heavy hydrogen) and tritium (superheavy hydrogen) fuse to make helium and a neutron – releasing 17.6 MeV of energy as fusion power. This is the easiest fusion reaction to initiate since it has a high reaction rate at low temperature (where "low" means 100–200 million kelvin). Tritium does not occur in nature as it decays with a short 12-year half-life to helium-3. Thus it must be "bred" from lithium using the neutron produced in the deuterium–tritium fusion reaction. Here, the neutron causes a tritium-breeding reaction with the isotope lithium-6, which comprises roughly 7.5% of naturally occurring lithium. The fuels for this fusion reaction are therefore deuterium and lithium, which are plentiful in seawater.

From dream to reality

The modern fusion programme really began in the closing moments of the Second World War at Los Alamos in the US, when Enrico Fermi and other members of the team that built the first atomic bombs speculated that a fusion reaction might be initiated in a plasma confined by a magnetic field. In May 1946 George Thomson and Moses Blackman of Imperial College London applied for a patent for a magnetically confined fusion device in which powerful magnets could be used to hold a plasma in place while it is heated to high temperatures.

By the early 1950s it was clear that the easiest fusion reaction to initiate is that of two isotopes of hydrogen – deuterium and tritium. To initiate significant fusion, a plasma of deuterium and tritium must be heated to temperatures of about 150 million kelvin. Some 10 times hotter than the centre of the Sun, this was a

At a Glance: Fusion energy

- Fusion power has the extraordinary promise of practically unlimited fuel, no carbon-dioxide production, good safety and insignificant land use
- Controlled fusion was realized in the 1990s by the Joint European Torus (JET) and the US Tokamak Fusion Test Reactor. JET needed more energy to run than it produced – 25 MW input power to the plasma produced 16 MW of fusion power
- We could reach net electricity production by building a reactor that can support the hot burning-plasma regime, where fast-moving fusion products self-heat the reaction, so that less input power is required
- Simulations and measurements predict that the ITER facility being built in France will reach this regime by having a less turbulent fusion plasma and a greater volume – therefore making more fusion and losing less energy – than its predecessors
- For commercial fusion, a wall and "blanket" for the reactor must be engineered that can withstand many years of heat and radiation without weakening

daunting goal. However, in 1997 scientists achieved it in a magnetically confined plasma at the Joint European Torus (JET) at the Culham Centre for Fusion Energy in the UK. JET produced 16 MW of fusion power while being driven by 25 MW of input power.

Eddington would no doubt be pleased with the scientific progress on his vision. But despite the successes, we are not yet at a point where we can generate commercial electricity and fusion's home stretch still involves significant challenges. Exactly what needs to be done to make a commercial fusion power source? What are the key scientific issues? How should countries position themselves to participate in a future fusion economy? These are essential questions. Before turning to them, however, it is worth addressing the most important question of all: why bother? Perhaps other energy sources would be simpler options. In reality, there are worryingly few long-term energy sources with sufficient resources to replace the roughly 80% of our energy that is generated by fossil fuels.

In the coming decades, current nuclear-fission technology will play a critical role in generating low-carbon electricity. But in the long term, aside from fusion, only solar and nuclear fission with uranium or thorium breeders (advanced reactors that breed nuclear fuel and so extend the resource of fission fuel) have the capability to replace fossil fuels. These technologies still need extensive research before they are ready to be deployed on a large scale. But despite this potential, it is clear, however, that no energy source offers the extraordinary promise of fusion: practically unlimited fuel; low waste; no carbon-dioxide production; attractive safety features and insignificant land use. These are compelling reasons to develop fusion even if success is not fully assured.

Self-heating fusion reactors

What then needs to be done to capitalize on JET's achievement of significant fusion power? The next stage is clearly to demonstrate that a plant producing a net amount of electricity can be constructed – something that JET was not designed to achieve. The ratio of fusion energy produced to the electrical energy consumed to initiate and sustain the reaction must be increased. This requires a self-heated plasma – one heated by the energetic helium nuclei produced in deuterium–tritium fusion (figure 1).

The National Ignition Facility (NIF) at the Lawrence Livermore National Laboratory uses a different approach to fusion than the magnetic-confinement method discussed here. The facility is designed to concentrate 500 TW of power onto a millimetre-scale fuel pellet using an array of 192 lasers. The fusion energy produced is expected to be roughly 10 to 20 times what the laser driver delivers as light. This would be a significant demonstration of fusion "burn", i.e. self-heating. However, the NIF laser is less than 1% efficient and thus the facility is still short of the critical demonstration that net energy production is possible.

For magnetically confined fusion, the crucial demonstration is at hand. Seven international partners – China, the European Union, Japan, South Korea, India, Russia and the US, together representing more than half the world's population – are now, after years

of delays, building a self-heated device called ITER at Cadarache in southern France (figure 2). Like JET, this experiment will have a magnetic configuration denoted by the Russian acronym "tokamak". ITER will be completed in 10 years and a few years after that is expected to be producing roughly 500 MW of output power from less than 50 MW of input power – a 10-fold amplification, or "gain", at least. One-fifth (roughly 100 MW) of the fusion power will be released as energetic helium nuclei, which get trapped by the magnetic field and self-heat the plasma. The target is to sustain this power level for a duration of 400 s or more. However, recent experiments using JET and other machines, coupled with detailed modelling, show that it should be possible to significantly increase that duration – and the gain. Even without these increases, ITER will generate industrial levels of fusion power while being largely self-heated; this is the *burning-plasma* regime. This demonstration of the scientific feasibility of high-gain fusion is a critical step on the road to fusion power.

But how do we know that ITER will reach these performance levels? The key physics parameter is the "energy confinement time", τ_E, which is the ratio of the energy in the plasma to the power supplied to heat the plasma, where the latter is both the self-heating due to the fusion-produced helium (one-fifth of the fusion power, $P_{fusion}/5$) and the external heating (P_{heat}). The energy confinement time parametrizes how well the magnetic field insulates the plasma – it might be thought of as roughly the time it takes the heat put in to the plasma to work its way back out. The plasma is sustained for many energy confinement times (in principle indefinitely) by the heating. Clearly, a larger τ_E makes a fusion reactor a better net source of power. The energy gain is defined as $Q = P_{fusion}/P_{heat}$. The deuterium–tritium fusion power produced per cubic metre of plasma at a given temperature and density (the fusion power density) can be calculated using the measured fusion cross-section (the reaction rate for a given fusion collision). In the temperature range 100×10^6–200×10^6 K, the fusion power density is approximately $0.08p^2$ MWm^{-3}, where the plasma pressure, p, is measured in atmospheres.

At high pressure the fusion power is large and the plasma is entirely self-heated ($P_{heat} = 0$ and $Q \to \infty$) – this is termed "ignition". Heating the plasma externally (supplying P_{heat}) reduces the net output and complicates the reactor design. Therefore, high gain is essential. The gain of a fusion device depends on the state of the plasma – specifically the fusion product, $p\tau_E$, and the plasma temperature, T. Ignition occurs roughly when $p\tau_E > 20$. In ITER the central plasma pressure will reach about 7 atmospheres and the confinement time is expected to be in the range 3.5–4 s (recall that ITER's plasma will be sustained for more than 400 s – perhaps thousands of seconds). A plot of $p_i\tau_E$ versus T enables a performance comparison for different tokamaks, where $p_i = p/2$ is the ion pressure in the centre of the toroidal plasma (figure 3).

Predictions of high power at ITER
The most challenging technical question faced by the fusion community is determining what the confinement time is and how we can be sure that it will reach 3.5–4 s.

2 ITER
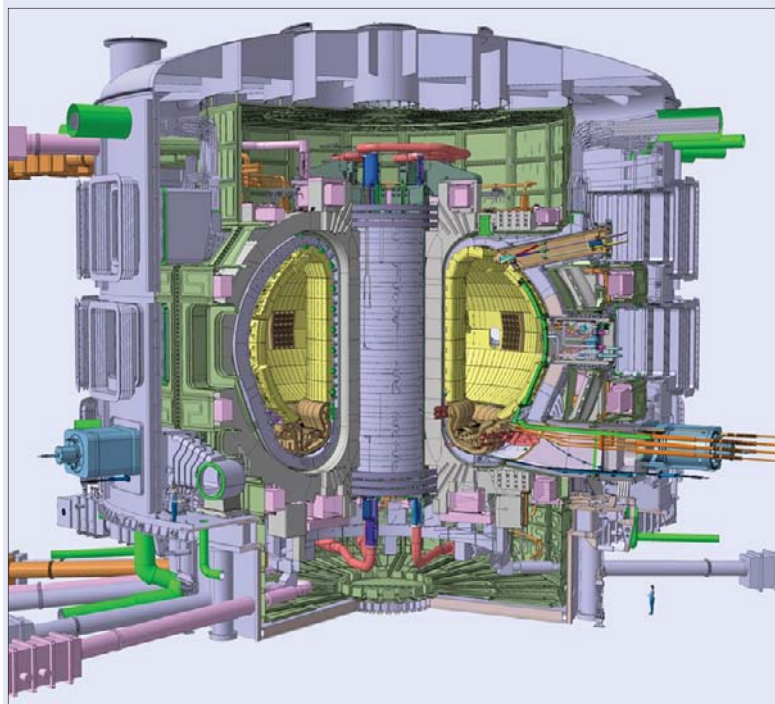
Now being built at Cadarache in southern France, ITER will contain roughly 830 m^3 of hot plasma inside a toroidal-shaped cavity. Confinement is provided by a magnetic field of approximately 5.2 T created by a niobium–tin superconducting coil at a temperature of 4 K. The plasma will be heated to fusion temperatures by radio waves and energetic neutral particles that are injected into the plasma. Once at fusion temperatures (about 200 million kelvin) ITER is expected to produce about 500 MW of fusion power for more than 400 s and be largely self-heated – such plasmas are termed burning plasmas. ITER is designed to have a "duty cycle" of at least 25% – i.e. the gap between burning-plasma shots is less than three times the shot duration.

We know that the loss of heat from magnetically confined plasmas is controlled by small-scale turbulence. The turbulence consists of plasma-density and electromagnetic-field fluctuations that cause little swirls of plasma flow – eddies. The turbulent fluctuations are essentially unstable sound waves driven by the temperature gradient in the plasma. Like convection in a saucepan, eddies transport hot plasma out and cold plasma in. Progress in tokamak performance over the last 40 years has been achieved by increasingly suppressing the turbulent convection of heat and thereby increasing τ_E. One of the scientific triumphs of the last decade has been the ability to calculate this turbulence using high-performance computers to provide state-of-the-art simulations (figure 4).

Detailed comparisons of the simulations and measurements show that in many cases the calculations are indeed correctly capturing the complex dynamics. There is, however, still room for improvement, especially in the intriguing cases where the simulated turbulence is almost entirely suppressed. The analytical theory of this turbulence is complicated and is only now just beginning to be understood. However, a qualitative understanding of the turbulent transport can be obtained from a simple random-walk argument based on the characteristics of the unstable sound waves that form the eddy structures. This argument yields the estimate $\tau_E \propto L^3B^2T^{-3/2}$, where L is the size

3 Progress towards the promised land
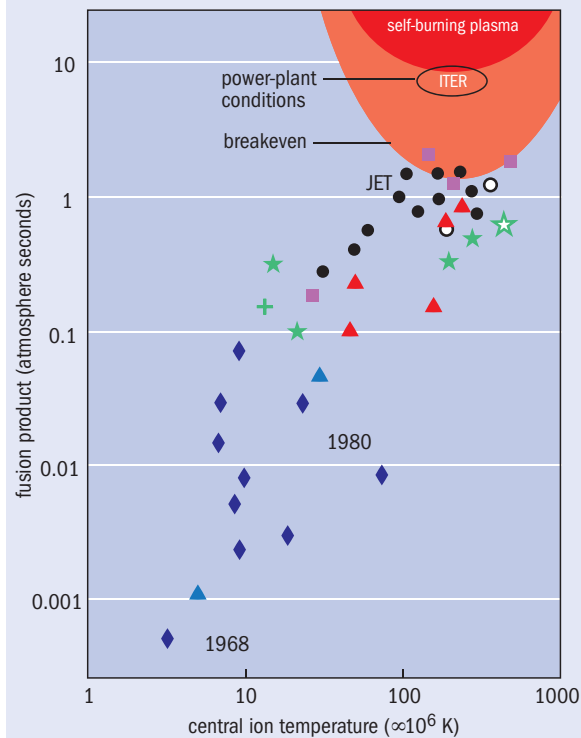
Selected data from different tokamaks demonstrate substantial progress over recent decades, with ion temperatures of more than 100 million kelvin now routine. With JET, an energy gain (Q) of about 0.7 has been reached – this is labelled as "breakeven" in this diagram. The Japanese experiment JT60 ran without tritium but if it had been using tritium, then the gain would have been 1.25. ITER is expected to obtain an energy gain of more than 10 – commercial reactors would need more than 20.

of the device, B is the magnetic field strength and T is its temperature. Clearly, bigger devices should perform much better due to the steep L^3 scaling. Indeed, empirical scaling derived from many experiments differs only a little from the simple estimate. ITER's energy confinement time has been predicted in two ways: first, by extrapolation from the existing machines using the empirical scaling; and second, using sophisticated local-transport models derived from simulations. These predictions are expected to be very accurate, with confinement times in the range 3.5–4 s. This prediction is the basis of our confidence that ITER will reach the self-heated burning-plasma regime. We can get a qualitative feel for the extrapolation using the simple random-walk scaling: JET achieves roughly $\tau_E \sim$ 0.5–1 s confinement times and therefore ITER (which will be roughly twice as big, 30% hotter and have a field approximately 30% larger) will have roughly $\tau_E \sim$ 4 s.

Blanket engineering

Given our current knowledge, it is more than reasonable to assume that ITER will achieve its goal of a burning plasma in the mid-2020s. However, as any engineer will confirm, there is much more to commercial power generation than simply proving a design is scientifically feasible. Indeed, several components of any future fusion reactor – in particular the systems

that breed tritium from lithium (the second reaction in figure 1) and convert neutron power to electrical power – have yet to be tested at any scale. The neutrons produced in deuterium–tritium reactions, which carry four-fifths of the fusion power, are not confined by the magnetic field and therefore leave the plasma and pass through the surrounding wall. Inside the wall there must be a complex system that absorbs the neutrons, extracts heat and "breeds" tritium from lithium – this is known as a "blanket".

There are many blanket designs but they all have a few things in common: they are typically 0.5–1 m thick, separated from the plasma by a steel wall and bounded on the outside by a steel shield. The blanket contains lithium, which absorbs neutrons from fusion to breed tritium (figure 1) that is then fed back into the plasma as fuel. Also in the blanket are neutron multipliers and a coolant used to flush out tritium and heat, which is used to power a turbine and generate electricity.

The blanket must satisfy some key requirements: to be economically viable it should operate robustly at high temperature in a harsh neutron environment for many years; and for tritium self-sufficiency it must breed more tritium than the fusion reactions consume. The technologies of the blanket, as well as the wall, are becoming a major focus of the fusion programme and will represent much of the intellectual property associated with commercial fusion. These reactor-system technologies are critical for a future fusion economy – we cannot wait for ITER's results in order to start developing them.

A prerequisite for a viable blanket–wall system is robust materials. Structural materials, breeder materials and high-heat-flux materials are needed. In typical reactor conditions the atoms in the first few centimetres of the wall facing the plasma will get moved, or displaced, by neutron bombardment more than 10 times per year. Each displacement causes the local structure of the solid wall to be rearranged. Often this will be benign but sometimes it can weaken the structure. Materials must therefore retain structural integrity in these very challenging conditions for several years. To minimize the environmental impact of fusion, the walls must also be made of elements that do not become long-lived radioactive waste following high-energy neutron bombardment.

We do not know for certain whether such materials exist, but several promising candidate materials have been proposed. For example, various special steels have been shown to have suitable structural properties in theoretical calculations and ion-beam tests undertaken at Culham and UK universities. But we will not know for sure until samples have been subjected to a fusion-type neutron-radiation environment. The International Fusion Materials Irradiation Facility (IFMIF) is an accelerator-driven neutron source being developed by the international research community to test small samples of the promising materials; its design team is based in Japan as part of the deal that brought ITER to Europe. The neutron spectrum of IFMIF will mimic the high-energy neutron spectrum of a fusion reactor. Samples will be irradiated in a beam of neutrons for several years to evaluate the changes in their structural properties.

Given our current knowledge, it is more than reasonable to assume that ITER will achieve its goal of a burning plasma in the mid-2020s

4 Heat loss through turbulence
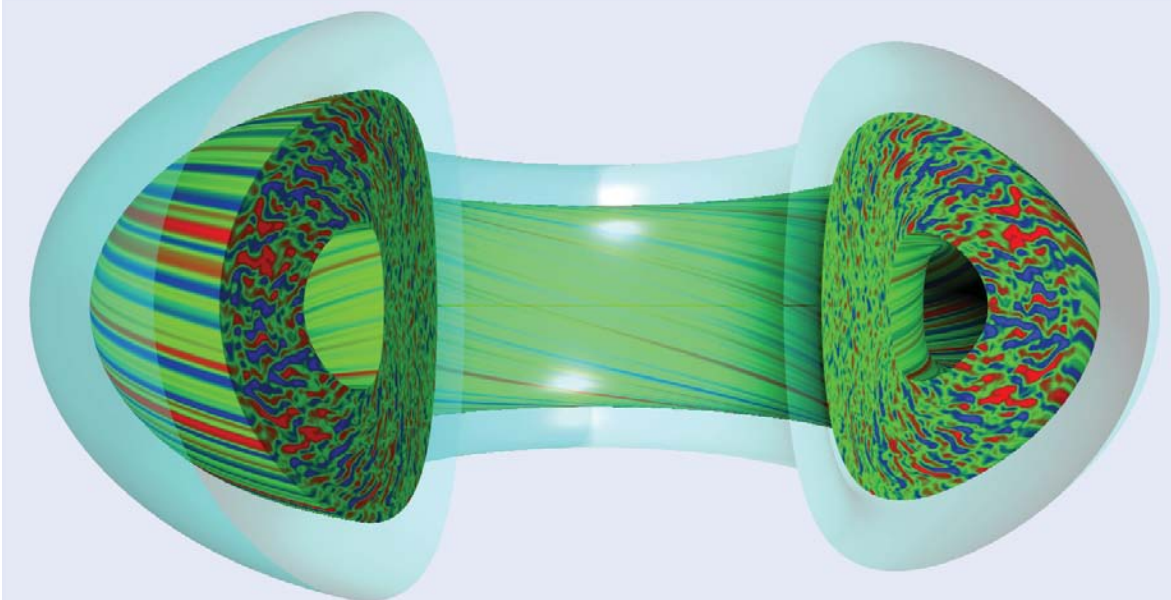
Jeff Candy and Ron Waltz

Fluctuations of plasma density caused by turbulence, as simulated for the DIII-D tokamak at General Atomics in La Jolla, California, using a computer code called GYRO. Magnetic field lines lie on nested doughnut-shaped surfaces – toroidal surfaces. In this image we can see two such surfaces, the turbulence between them and two cuts across the surfaces. The hot middle of the plasma is omitted. The field lines are not shown but the fluctuations are elongated along the magnetic field lines and are thus visible as the red and blue streaks along the toroidal surface. Turbulent flow is roughly along lines of constant colour and is perpendicular to the magnetic field lines. As can be seen from the cuts across the surfaces, the swirls – eddies – are shorter in scale across the field (they are a few times the width of the helical orbit of the ions about the field lines). GYRO solves kinetic equations for the rings of charge formed by the helical motion of particles around the magnetic field lines. The fields are calculated from Maxwell's equations using the calculated charge and current.

We need a testing facility

If, as expected, ITER proves to be successful, then blanket development is probably the critical path for fusion. Blanket designs are being developed and tested with weak sources of neutrons, and it appears that these designs will breed tritium efficiently enough to be self-sufficient. But they must be tested at full neutron power before we can ensure a reliable commercially viable system. Although test-blanket modules will be placed in the walls of ITER in the later stages of operation, definitive tests require a continuous neutron flux of 1–2 MW m^{-2} for several years, which will not be technically possible at ITER. Thus I believe that a "component test facility" (CTF) that can deliver reactor-level neutron flux over many square metres is needed to significantly accelerate the development of blanket and wall structures. For such a device to be affordable it must be compact with low power consumption.

Researchers at Culham have pioneered a compact device called the spherical tokamak that is a prime candidate for a CTF. Indeed MAST (the MegaAmp Spherical Tokamak) has achieved impressive plasma conditions at a very modest scale. Calculations and measurements suggest that MAST achieves good confinement by suppressing the turbulence by spinning the plasma at supersonic speeds. The National Spherical Tokamak Experiment (NSTX) at Princeton in the US also operates at about the MAST scale.

Results from these devices suggest that the spherical tokamak is an ideal candidate for a compact and affordable fusion device – i.e. a suitable candidate for a CTF. Culham and the Oak Ridge National Laboratory in the US have therefore developed conceptual designs of CTFs based on spherical tokamaks. These facilities could test whole components of the blanket and wall at full power for many years. Both Princeton and Culham are upgrading their machines to prove the viability of these conceptual designs. The MAST upgrade will deliver near-fusion conditions, sustained plasmas and a test of the new exhaust system for gaseous plasma-burn products – the Super-X divertor.

If the MAST upgrade confirms the viability of a spherical CTF, then one could be built during the early years of ITER's operation. Wall and blanket development on the CTF coupled with ITER's programme could enable the construction of the first demonstration reactors in the 2030s. The current international programme has no plans to build a CTF – but surely it is essential if we are to deliver commercial fusion when it is needed.

It seems inevitable, given what has been achieved, that Eddington's dream will come true eventually – but when? Although we cannot say for sure, for a world that is hungry for energy, a reduction of the time to commercial fusion by even one decade could have an enormous impact. ∎

More about: Fusion energy

ITER: www.iter.org
NIF: https://lasers.llnl.gov
K Ikeda *et al.* 2007 Progress in the ITER physics basis *Nuclear Fusion* **6** 47
R Pitts, R Buttery and S Pinches 2006 Fusion: the way ahead *Physics World* March pp20–26