Module 5

Deployment

Authored by: Bas Marechal, Sander Korteweg
and Jens Peter Frankemölle

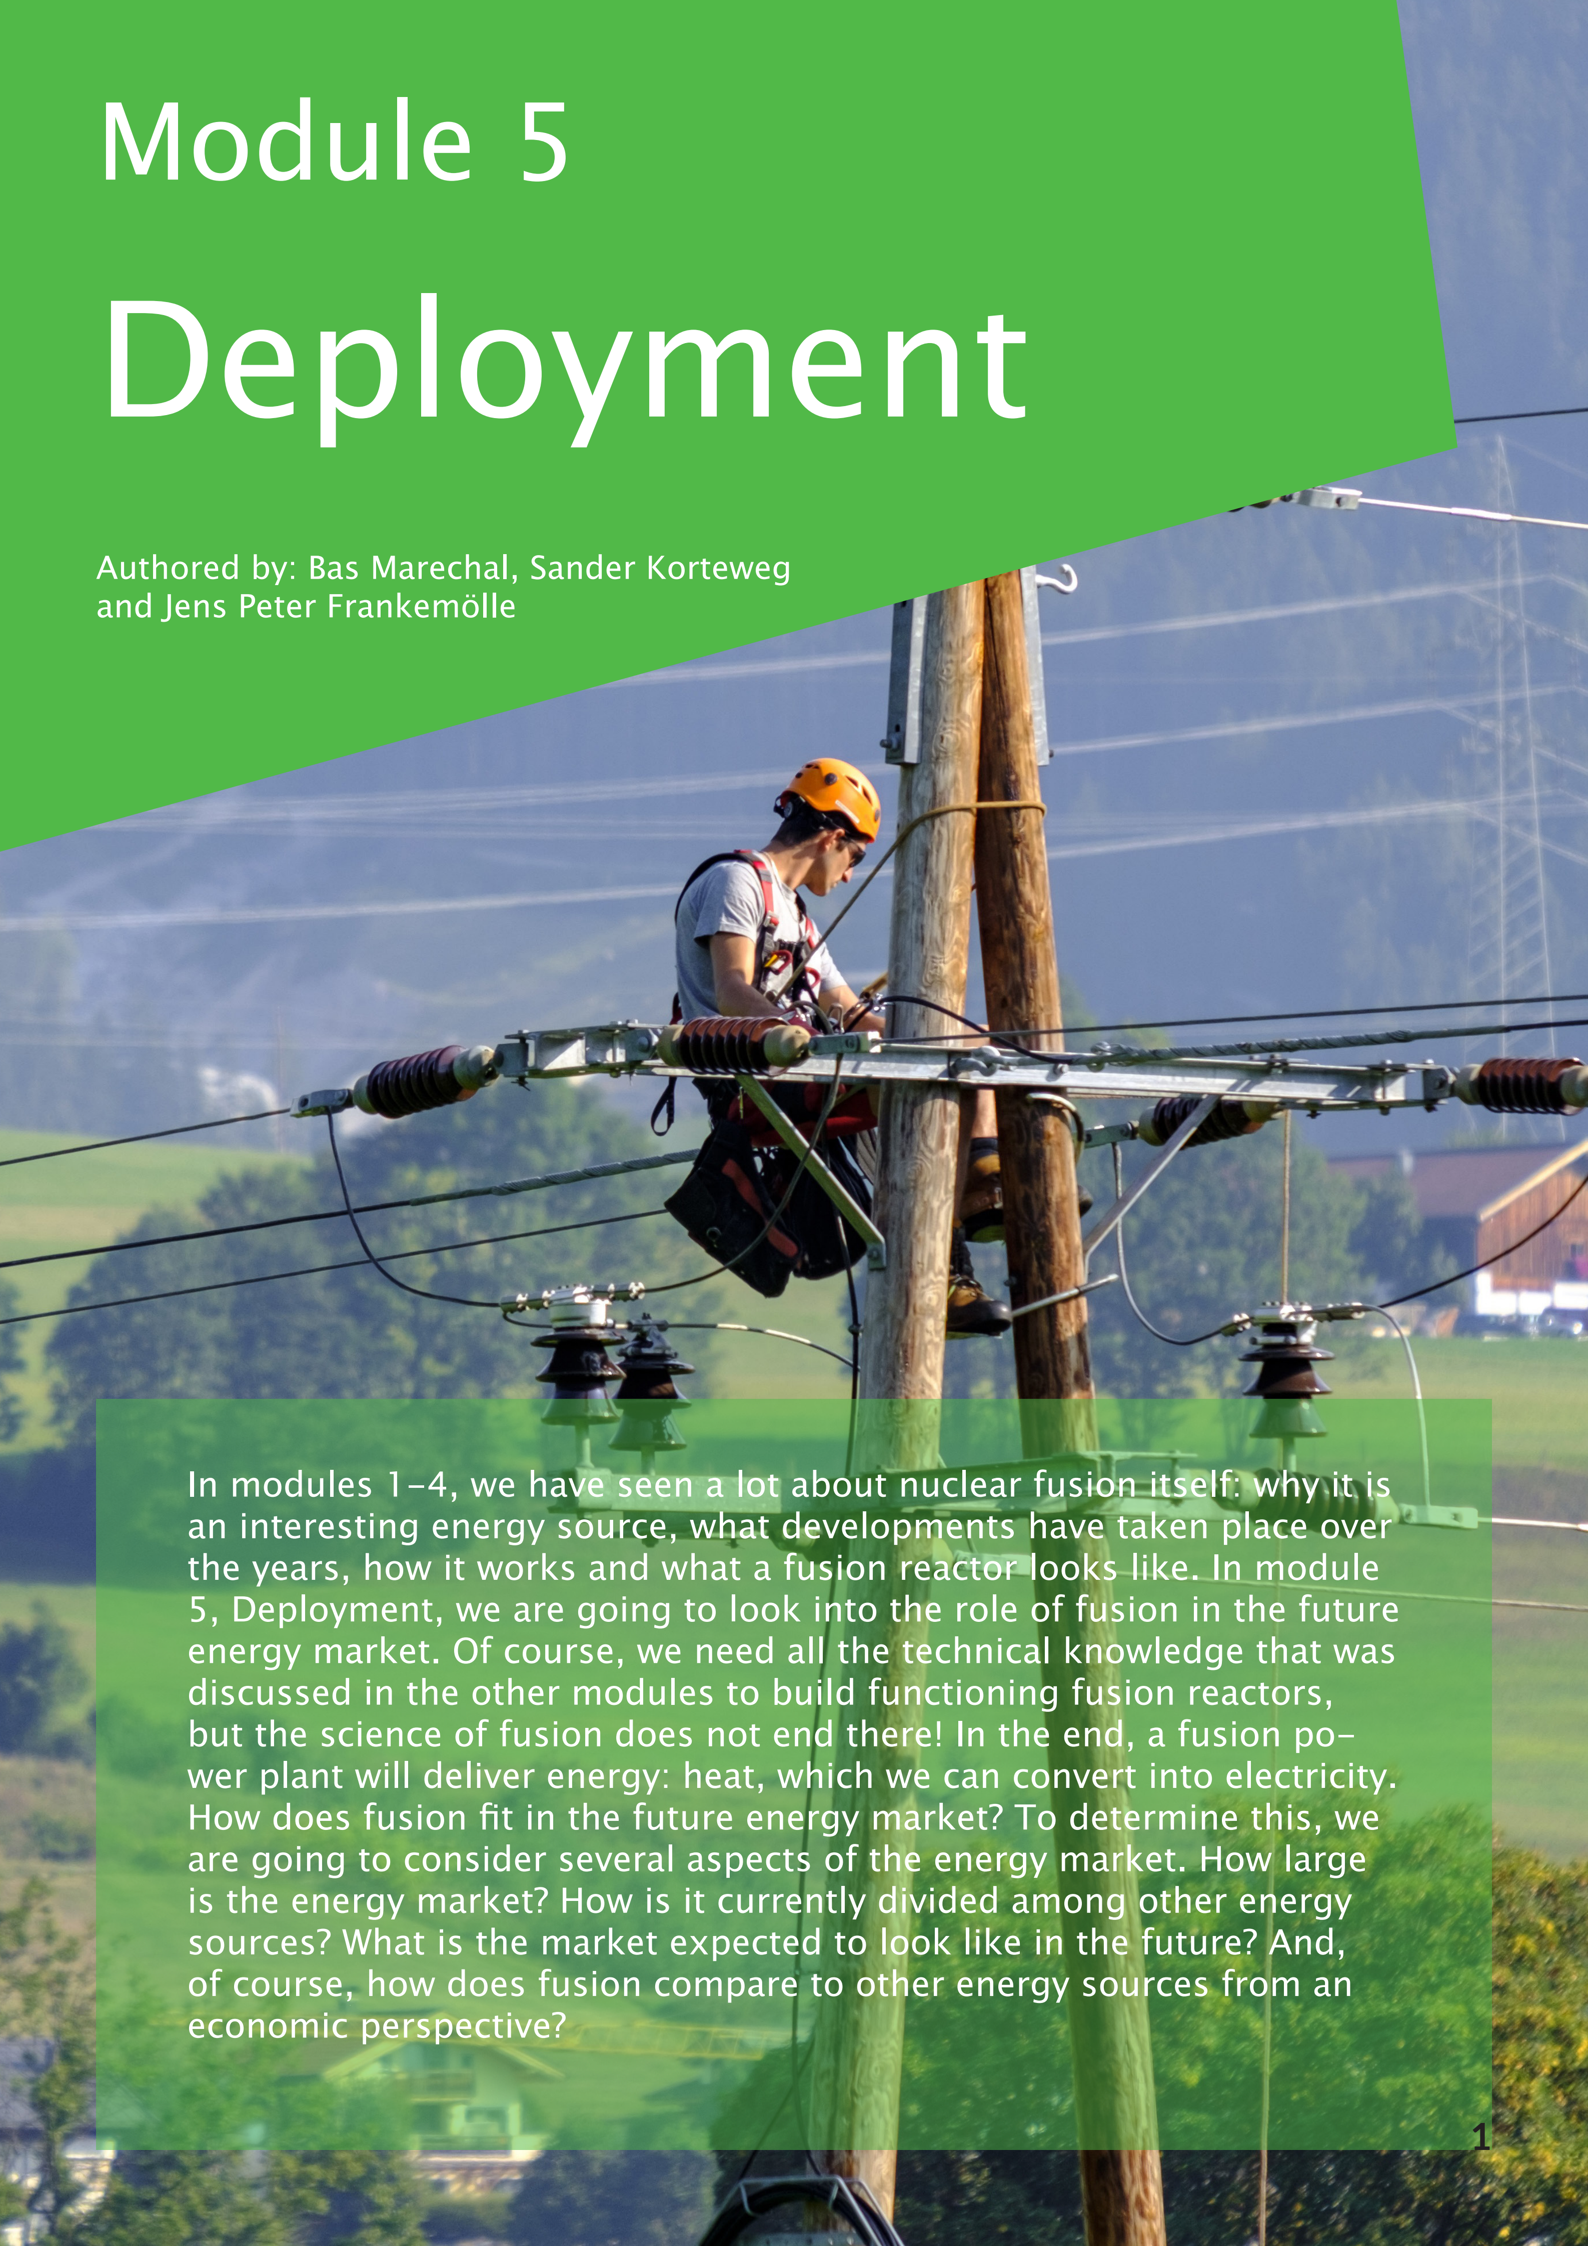

In modules 1-4, we have seen a lot about nuclear fusion itself: why it is an interesting energy source, what developments have taken place over the years, how it works and what a fusion reactor looks like. In module 5, Deployment, we are going to look into the role of fusion in the future energy market. Of course, we need all the technical knowledge that was discussed in the other modules to build functioning fusion reactors, but the science of fusion does not end there! In the end, a fusion power plant will deliver energy: heat, which we can convert into electricity. How does fusion fit in the future energy market? To determine this, we are going to consider several aspects of the energy market. How large is the energy market? How is it currently divided among other energy sources? What is the market expected to look like in the future? And, of course, how does fusion compare to other energy sources from an economic perspective?

5.1 Fusion power

The global energy market is constantly changing and for all manner of reasons. To name a few: economies are evolving, new technologies are emerging and regulations to combat climate change are intensifying. And so the way that we **consume energy** changes.

Much the same can be said of the way that we **produce energy**. The sources that we draw from are changing. Power plants that run on fossil fuels use carbon capture to reduce emissions. Slowly but steadily, solar panels and wind turbines are becoming more efficient. And in the early 20s of this century, the taboo on nuclear energy generation seems to be lifting – albeit slowly.

If we want to understand the chances that nuclear fusion has to actually be incorporated in our future energy mix, we need to understand how the energy market works and how it is changing. To do so, we first need to learn a new vocabulary. Words like *overnight costs*, *baseload* and *levelized cost of electricity* will feature plenty in this module.

Our starting point is the fusion reactor from a power plant perspective. How much does it produce? And how does that stack up against other production methods? Before we go there, however, we will first need to set the record straight on an error that is commonly made when talking about energy and electricity.

Energy versus electricity

In the media, energy and electricity are often used interchangeably. However, they are not the same. Electricity is a form of energy, but there are more forms of energy, such as heat or motion. In 2019, electricity usage only amounted to 19.7% of the total energy consumption. At first glance, this seems a startlingly low percentage but on closer inspection it makes sense. Looking at your own home situation, any of the following might apply: the car sitting in the garage is runs on petrol or diesel, your house is heated by gas and so is your food, because the stove runs on gas as well. All of that counts towards your energy consumption, although none of it runs on electricity. 80% of all energy that is used worldwide is non-electric.

Looking at energy consumption at a global scale, we find that those same categories that we found in our homes – transportation and heating – account for a significant chunk of total consumption. A third category, industrial processes, also accounts for a significant fraction. Eliminating fossil fuels from the energy market thus requires more than 'green electricity' alone.

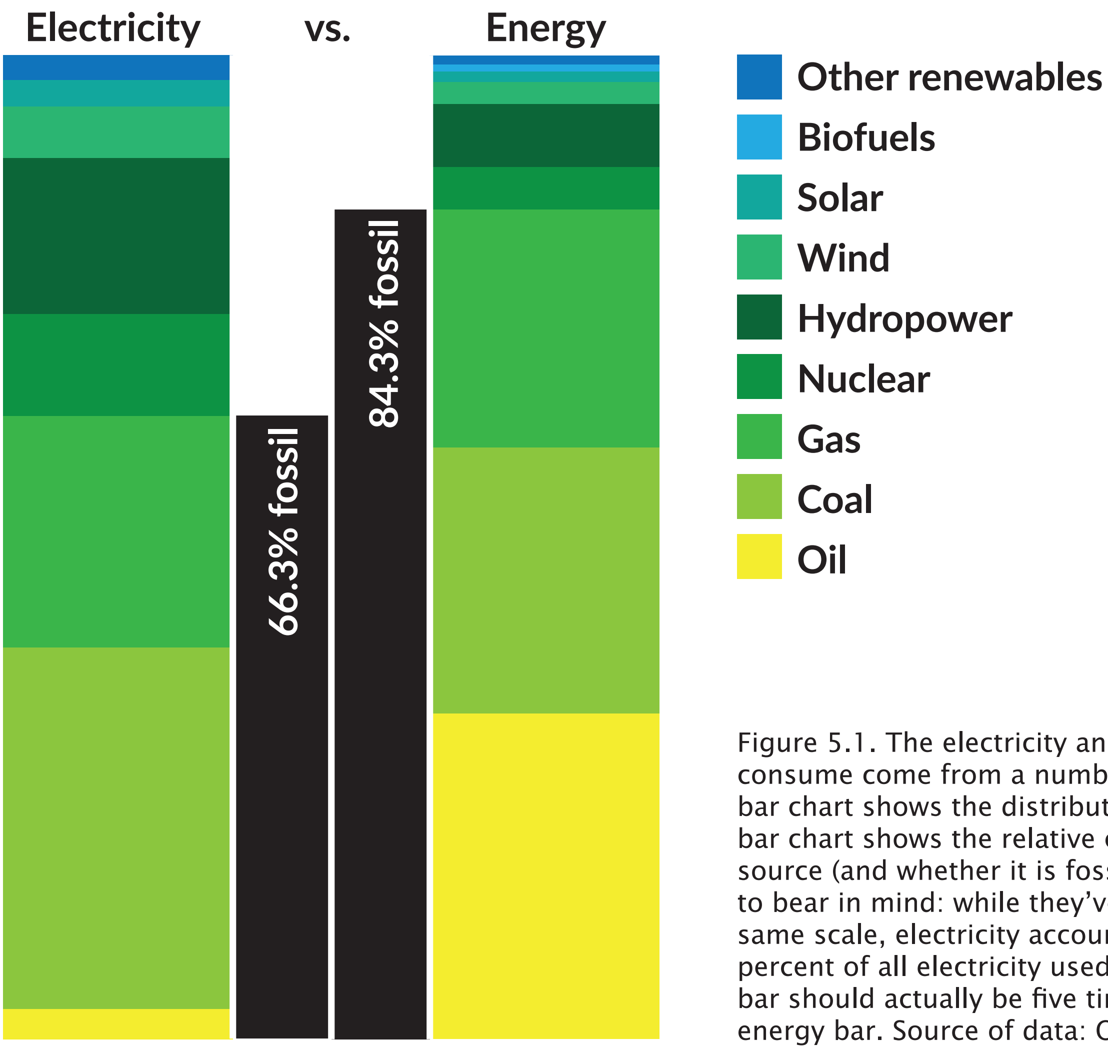

Figure 5.1. The electricity and energy that we consume come from a number of sources. This bar chart shows the distribution in 2019. This bar chart shows the relative contributions of each source (and whether it is fossil or not). On thing to bear in mind: while they've been drawn at the same scale, electricity accounts for only 19.7% percent of all electricity used and so the electricity bar should actually be five times smaller than the energy bar. Source of data: Our World in Data.

Sources of electricity and energy

Figure 5.1 on the previous page and table 5.1 on this page show how the global electricity consumption and global energy consumption were divided over different sources in 2019. Notice that nuclear fusion is not part of the market yet, since the technology is still being developed, but that nuclear fission is.

From figure 5.1, we can immediately see that fossil fuels are still dominating: they accounted for 84.3% of global energy and when looking purely at electricity consumption, we see that still 63.3% of all electricity came from fossil fuels in 2019. Even though electricity is increasingly generated through renewables, this still only accounts for 19.7% of all energy in the world.

All in all, it shows that the resources used to produce electricity are different from those used to produce energy. A journalist that does not understand the difference between electricity and energy might erroneously write: “a third of all energy is produced from green sources”, while he or she should have written: “15% of all energy is produced from green sources.” The first is considerably more positive than the second, but very much wrong.

Type	Electricity	Energy
Other renewables	2.5%	0.9%
Biofuels	<0.1%	0.7%
Solar	2.7%	1.1%
Wind	5.3%	2.2%
Hydropower	15.8%	6.4%
Nuclear	10.4%	4.3%
Gas	23.5%	24.2%
Coal	36.7%	27.0%
Oil	3.1%	33.1%

Table 5.1. Global electricity consumption and global energy consumption per source in percentages. These are the same as in figure 5.1. Source: Our World in Data.

Fusion electricity production

Okay, so we have now learnt that energy and electricity are not the same thing. In this module, we choose to focus mainly on the **electricity market** to limit the length of this module somewhat and because it is what a fusion reactor is set to produce, after all. And that begs the question: how much electricity is a fusion reactor going to produce?

Well, the goal for the first fusion power plants is to deliver 1 gigawatt (GW) of electric power to the grid. If you multiply this 1 GW by the total number of hours in a year (8 760 h) you get the total amount of energy generated: 8 760 GWh. However, even the best power plants need maintenance sometimes so the reactor cannot produce electricity year-round. This is captured by the plant **availability**. If we assume a plant availability of 95%, we find that the total amount of energy generated per year is actually $0.95 \times 8\,760 = 8\,330$ GWh.

Classroom Exercise 5.1

The unit of kilo-watt-hour (kWh) provides a way of measuring large amounts of energy. It is the energy that is used if we let a machine with a power of $P = 1$ kW run for a time of $t = 1$ h. The kWh is not an SI unit. We can convert it to Joule (J), the SI unit of energy, by expressing the power in watt and the time in seconds. We then find that

$$1 \text{ kWh} = 3\,600\,000 \text{ J}$$

(a) Calculate this conversion rate yourself.

(b) Using figure 5.2, calculate a household's yearly electricity consumption in J.

(c) Calculate how many households can be powered in a year by a 1-GW nuclear power plant.

So how does this 8 330 GWh compare to the total electricity consumption on Earth? And to that of a big city? Or of your house? You'll find these numbers in figure 5.2. And how does the energy generated by nuclear fusion stack up to other modern-day sources of electricity? Wind mills, solar panels and even a typical coal power plant produce less energy than our nuclear fusion reactor. See for yourself in figure 5.3.

Using these numbers we can calculate that one fusion power plant has the potential to meet the electricity demand for roughly 2.4 million households. And if we want to provide all electricity by fusion power plants, more than 2 500 power plants would be necessary. The large numbers of plants that we need may seem strange at first, if we look at the number of households that just one plant can serve. However, households are not the only driver for electricity demand: the industry and (to a lesser extent) the transport sector are large consumers as well.

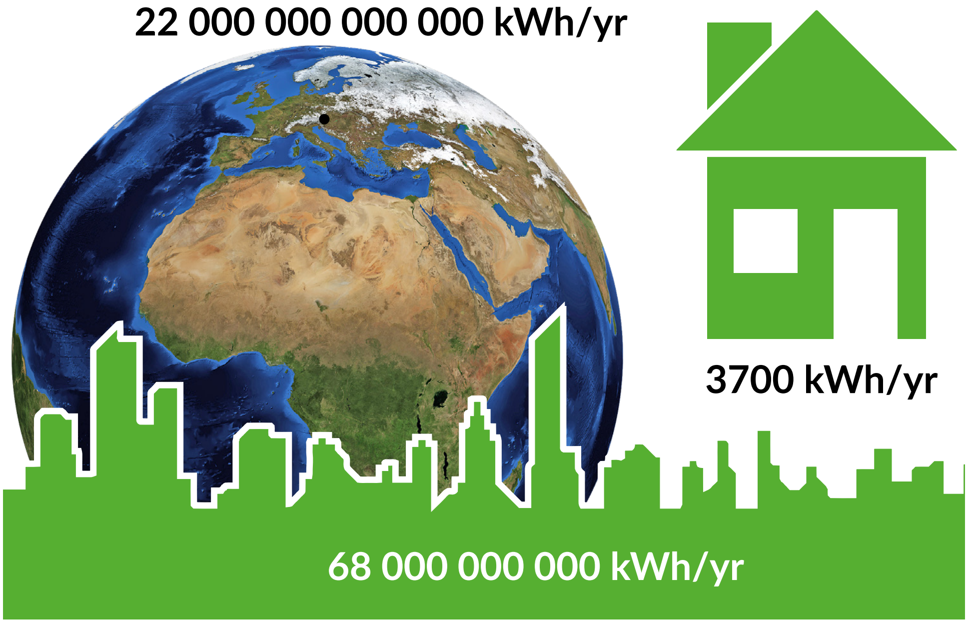

Figure 5.2. Comparison of electricity consumption of a household, a city and the world. Data: ODYSSEE-MURE, Our World in Data and University of Ontario.

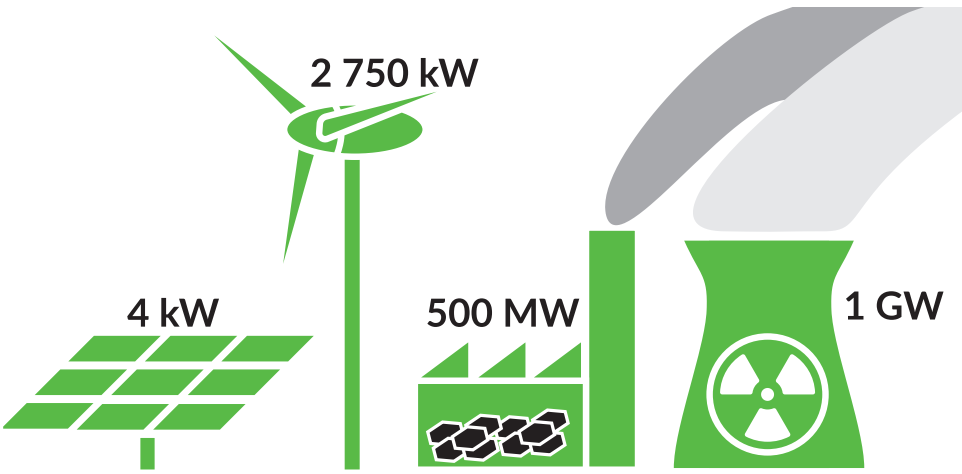

Figure 5.3. Comparison of electricity produced from a fusion power plant to the consumption (in 2018) on different scales. Data: energy.gov and McGinley.

ASIDE

Producing electricity

One of the most important laws in physics is the law of conservation of energy. It states that energy can never be created or destroyed, but can only be converted into different forms and states. When we speak of "producing" or "generating" energy, we do not create energy; we convert energy from a not-so-useful form into a (more) useful form.

For example, almost all power plants work on the same principle as the steam engines of the past: chemical bonds of the fuel (e.g. oil/gasoline/coal) are broken and rearranged to generate a certain amount of heat, used to expand, move or turn something (e.g. a piston or turbine). From this newly generated motion we can generate electricity by a dynamo: a moving magnet inside a coil that induces a current, resulting in electricity. In this way, energy is converted from a not-so-useful form (e.g. a lump of coal) into a useful form (e.g. electricity).

Classroom Exercise 5.2
So far, we assumed that a fusion reactor could run for 95% of the time. However, this is highly optimistic. Additionally, the power plant might not be able to always run at 100% when it is in use. Therefore, the concept of capacity factor was introduced. The capacity factor is the time-averaged power of a power plant divided by the power the reactor could deliver if it would run at maximum power for that same period.

If a power plant were to always run at full capacity (i.e. 100%), then the capacity factor would be 1.0. If it runs below its optimal capacity, the factor is between 0.0 and 1.0. The capacity factor for the first generation of fusion reactors is expected to be about 0.6.

(a) Calculate the amount of electricity produced in a year.

(b) Calculate how many households can be provided for with a single plant.

(c) Calculate how many plants we would need to provide for the electricity demand of the entire world.

5.2 Changing demand and supply

Since we still need to wait for about forty years until the first fusion power plants see the light of day, it is interesting to not only look at the current energy and electricity markets. What will the world look like in, say, 2040? While 2040 is quite a bit earlier than the 2060s in which the first nuclear fusion plants can be expected, this comparison will teach us about the way that the energy landscape is changing.

Increasing electricity demand

We saw in module 1 that the global energy demand is still increasing. To get a sense for this increase, we can look at figure 5.4, which shows the global demand for different energy sources (including electricity) in 2018 and the expected demand in 2040.

In figure 5.4, we can see that energy demand is expected to increase by almost 25% in only a time span of 20 years, which is mostly the result of increasing demand for electricity and oil. It is also clear that the electricity demand is expected to grow the fastest in both an absolute and relative sense.

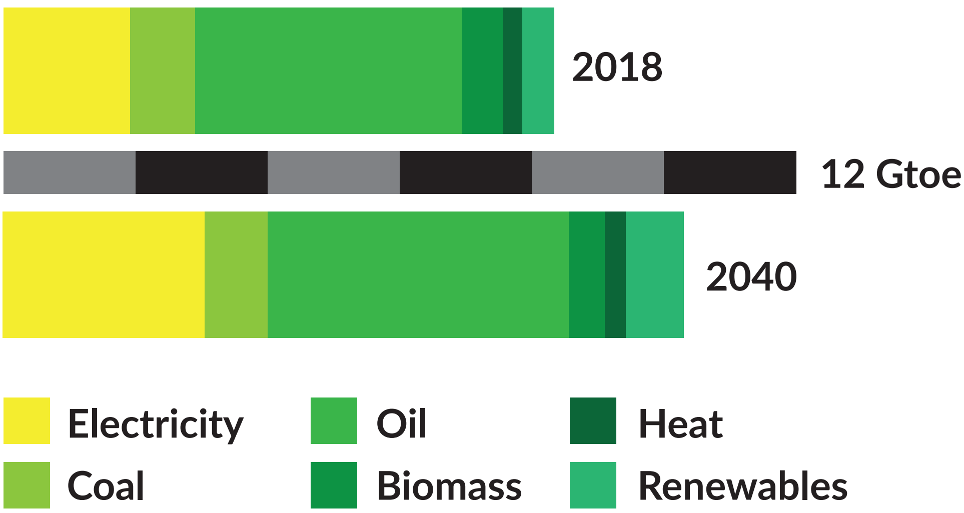

Figure 5.4. Absolute energy consumption in 2018 compared to absolute energy consumption in 2040. One Gtoe is one gigatonne of oil equivalent and thus equal to 1,000,000,000 toe. We see that the electricity demand will have grown by about 2 Gtoe in 2040. Data: WEO, 2019.

One explanation for the increasing total energy demand is the population growth we saw in module 1. Another reason for this is that developing countries (such as Brazil and China) are expected to consume more energy in the future to stimulate growth of their economies. Both of these causes continue well beyond 2050. So even if the goals of the Paris agreement are achieved, we would still need to apply large efforts to keep our energy sustainable.

However, population growth and the increasing development of countries cannot explain why electricity demand grows faster than demand for other energy sources. The reason behind the relative growth of electricity is a matter of policy. All sectors are forced to switch to sustainable energy. Because renewables, like solar and wind, tend to produce electricity, most sectors focus on moving from fossil fuels to electricity. As an example, think of the increasing number of electric cars. The growth in electricity demand means that there will be a larger market to provide for. The addition of a new energy source that produces electricity could be very helpful.

Classroom Exercise 5.3
In figure 5, the energy consumption per fuel is given in terms of ton of oil equivalent (toe). This unit is often used to express very large amounts of energy. However, it is not an SI unit. For this, we need to convert it into joules (J). Assume that

1 toe = 11.63 MWh

(a) Convert the electricity consumption and total energy consumption in both 2018 and 2040 to kWh.

(b) Convert the electricity consumption and total energy consumption in both 2018 and 2040 to J.

Electricity shifts towards renewables

Since our main focus is still with the electricity market, let us focus on the yellow section of figure 5.4. What means of production make up the electricity mix nowadays, and what shift do we expect in the next 20 years?

Figure 5.4 showed the changing energy demand by fuel, and it was clear that electricity demand will have increased by 2040. Figure 5.5 now shows the projected electricity suppliers in 2040.

This figure already shows an interesting projection the increase in demand in clean sources, mainly wind and solar PV. Meanwhile, the share of fossil fuels has decreased to 48%. Don't be fooled by these numbers, however: the supply from all sources, except for oil, does increase in an absolute sense due to the strong increase in demand. Therefore, the shift we see indicates that we use clean energy sources as much as we can to increase new demand, but that it is not yet sufficient to replace fossil fuels.

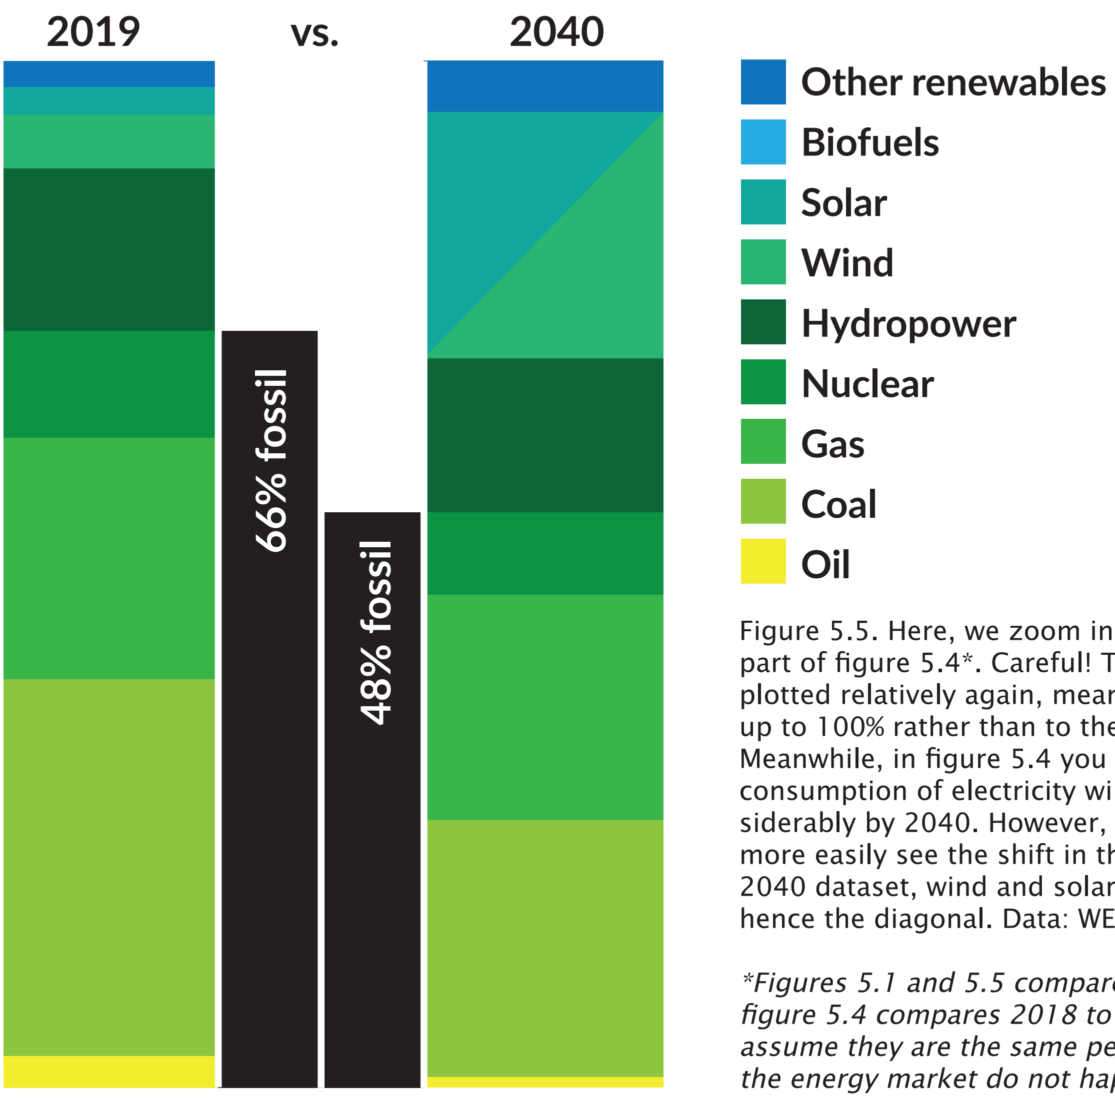

Figure 5.5. Here, we zoom in on the 'electricity' part of figure 5.4*. Careful! These contribution are plotted relatively again, meaning that the bars add up to 100% rather than to the total consumption. Meanwhile, in figure 5.4 you can see that the total consumption of electricity will have increased considerably by 2040. However, in this way we can more easily see the shift in the distribution. In the 2040 dataset, wind and solar were aggregated, hence the diagonal. Data: WEO, 2019.

**Figures 5.1 and 5.5 compare 2019 to 2040, while figure 5.4 compares 2018 to 2040. However, we assume they are the same period (major shifts in the energy market do not happen in a year).*

Baseload becomes critical...

Eliminating fossil fuels from the electricity market creates a void in supply, which needs to be filled up. This void in the electricity generation will come in the form of a so-called **baseload of electricity**, which means that it has to be delivered continuously during the day. Fossil fuels are excellent at delivering a baseload, since as long as you have fuel, the power plant can run. This is relevant for the production of a minimum level of electricity over time.

On the other hand, most renewables have fluctuating energy production and tend to be unreliable for baseload electricity: seasonal changes lead to less overall energy production during times in which most energy is needed (wind turbines and solar cells generate less energy in winter, when we need more light and heat).

Nuclear fusion is an electricity source that may provide a baseload supply in the future.

...or the entire grid gets reworked

The reason that we require baseload energy is that our current grid was built with a lot of baseload sources in mind. If we want (or are forced) to move away from baseload, that requires us to rework the entire grid.

We will compare fusion to its competition later in this chapter, but one option is worth discussing here: it is still debatable whether continuous electricity production (baseload) will remain necessary at such a large scale. Instead, we could also make large changes to the electricity grid. This could be done mainly in two ways.

Using a **super grid** – with a super grid, electricity can be produced on a large scale in remote areas and then be transported to urban areas. For example, solar panels can generate more electricity in a desert. This electricity can then be stored and transported to cities across different continents.

Using a **smart grid** – within a smart grid, there is a lot of active communication to move electricity around. In this way, electricity can be transported from places with an excess to places with a shortage. For example, if it is a sunny day, you may be generating more electricity with solar panels than you need to keep your drinks cool, so the grid can take some of your excess and use it for your neighbour that does not have solar panels yet.

Both of these options would reduce the need for continuous electricity production: we can either store the electricity that we produce with **intermittent** (fluctuating) sources, or transport excess electricity around the grid. This way, wind and solar energy can grow significantly on the electricity market, and we still have to consider them as competitors for fusion. There is, however, one drawback of these grids: they still need to be developed, which is expensive, and the storage capacity we can achieve is currently limited. Unfortunately, as we will see, the same holds for nuclear fusion, so we cannot simply say that either option is superior.

Summarising

So, electricity demand will be on the rise in the coming decades. Meanwhile, many of our current sources of electricity will no longer be available. Coal and gas and oil are a no-go if we are serious about combatting climate change.

Solar PV and off- and on-shore wind turbines will be a key part of the future mix of electricity sources. However, as these types of electricity production are inherently intermittent, we are faced with a serious issue. Our current grid is built with many baseload (continuous supply) electricity sources in mind, and we cannot just swap them with renewables. So, we need to rework our entire grid with renewables in mind or we need to develop alternative baseload technologies like nuclear fusion.

Baseload is often presented as a big advantage of nuclear fusion, and we now know why. However, we have also seen that there are alternative strategies in which baseload is less important. Much boils down to the choices we make and the alternatives we have. In the next section, we will focus on a couple of other possible advantages of nuclear fusion.

5.3 What can fusion add to the mix?

In the previous chapter we have learned that the 2040 electricity market will be quite a bit different from the one that we know today. Imagine the contrast between today and, say, the 2060s. As a consequence, it is quite difficult to answer the titular question: “What can fusion add to the mix?” In this section, we cover a few advantages of nuclear fusion that will likely still be relevant in the 2060s (and perhaps then more than ever).

Baseload electricity

The first one that we want to mention will not come as a surprise if you read the previous chapter: baseload. Classical downsides of solar and wind energy are that the Sun is not always shining and that the wind is not always blowing. Solar and wind are **intermittent** electricity sources.

Now, your immediate thought is probably: “Right, the Sun doesn’t shine at night!” However, that day-night cycle is actually not the main problem. It is the seasonality that really gets us. On average, there is less wind and less sunlight in the winter. That is a big issue, because as you can imagine, winter is the time of year that we need most energy. This is quite a setback compared to current-day electricity sources. Fossil fuels can provide electricity year-round at the flip of a switch, and we’re trying to replace them by sources that cannot. This is the baseload problem that you are already familiar with.

Nuclear fusion is a baseload electricity source. We’ve learned quite a bit about the way that nuclear fusion works in previous module and put at its most basic it works no different from a current-day coal plant. Put in fuel, flip a switch and voilà, you have power. There is one big ‘but’. For a useful baseload source it is important that the output power can be varied to match the intermittent sources: e.g. if the Sun is shining extra brightly, you want to be able to lower the output of your fusion reactor. For fusion to be a good baseload source, it needs to be able to adjust its output power or to have an easy on and off switch, which is not yet a given.

Land use

Several countries in Europe are already struggling to allocate sufficient space to all the different types of activities that occur in a country: housing, agriculture, industry and nature all compete for the same land. If you build a huge solar park on a stretch of empty land, that land is no longer available for farmers to farm on or for residents to live on.

The amount of electricity that can be produced per unit area differs greatly from technology to technology. One big reason that fossil fuels are so popular is that they pack quite a punch. There is a lot of energy contained within a single kilogram of fuel, and as such you only need to burn comparatively little of it (and thus need comparitively little land). The same cannot be said for solar and wind power, however, which is another downside of transitioning from fossil fuels to these types of renewables.

Nuclear fusion has by far the largest energy released per unit mass of all the energies known to man. The reason for this is that the burning of fossil fuels is a chemical process while the burning (rather: fusing) of deuterium and tritium is nuclear process. The energy contained in nuclear bonds is much larger than the energy contained in chemical bonds, such that the energy released in a reaction is much larger. This advantage is offset somewhat by the fact that so many auxiliary installations are required. There's far more than just the tokamak and all those installations need a piece of land. Still, nuclear fusion does quite well as can be seen from figure 5.6.

Classroom Exercise 5.4

There's quite a few auxiliary installations required besides the tokamak itself. What installations can you imagine (or remember)? Why are they needed?

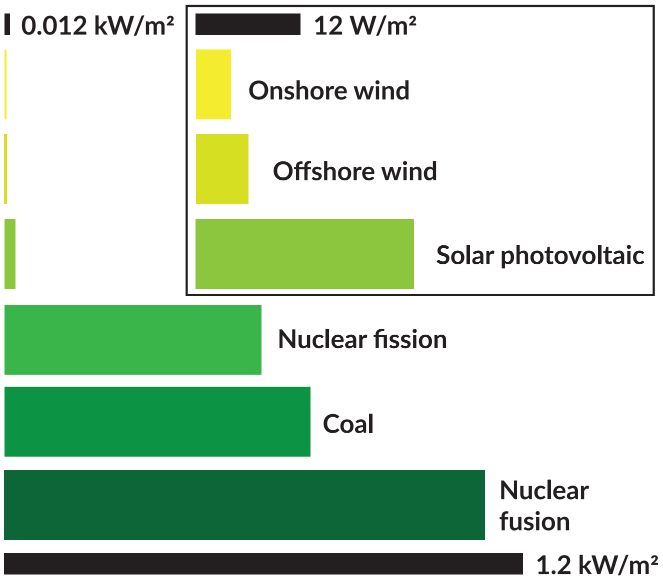

Figure 5.6. Bar graph comparing the energy densities per unit area for different energy sources. Data: MacKay, 2008 and Hore-Lacy, 2011.

Fuel depletion

This is quite a straightforward one. Aside from emitting greenhouse gases, fossil fuels come with another major downside. There's a limited supply of them. Some researchers think we may run out of oil and gas around the year 2045 and out of coal around 2115. Meanwhile, in 2019 two thirds of the global electricity production came from fossil fuels. We previously saw this in figure 5.5.

Like wind and solar energy, nuclear fusion does not suffer from fuel depletion. At least not anytime soon, as it is expected that there is enough fuel for at least a 1000 years to come. We will see in the last chapter of this module that fusion fuel comes with its own share of problems, but - if we get everything to work as envisaged - fuel depletion should not be one of them.

Meltdowns, weapons and waste

Finally, there are three 'classical' arguments against nuclear energy in general: **meltdowns**, nuclear weapons (**proliferation**) and long-lived **radioactive waste**.

The most infamous example of a meltdown is for sure the accident at Chernobyl in Ukraine in 1986. A meltdown is a runaway nuclear reaction that can result in the reactor overheating and melting down. Modern-day reactors are incredibly well-protected against such a meltdown occuring, but no safety system is perfect as

ASIDE

Levelized Cost of Electricity

And what if we want to look at just the money? Which source of electricity is the cheapest? Comparing the costs of different electricity sources is not as simple as you might think.

To compare the economics of different electricity sources, we can look at the **Levelized Cost of Electricity (LCOE)**. The LCOE is the total cost of a power plant throughout its lifetime (so construction, maintenance, fuel, etc.) divided by the amount of electricity it produces (accounting for some interest). This can be seen as the minimum price for which the electricity needs to be sold to make a profit.

Figure 5.7 shows the expected LCOE for different energy sources for the year 2050. Please note that the cost for fusion electricity will be higher than mentioned here when it first hits the market, because making the first reactors is always more expensive than making the hundredth reactor. From figure 5.7 we see that, in the future, fusion is estimated to be about as expensive as onshore wind. From these LCOE estimates, it seems that fusion could compete in the future electricity market.

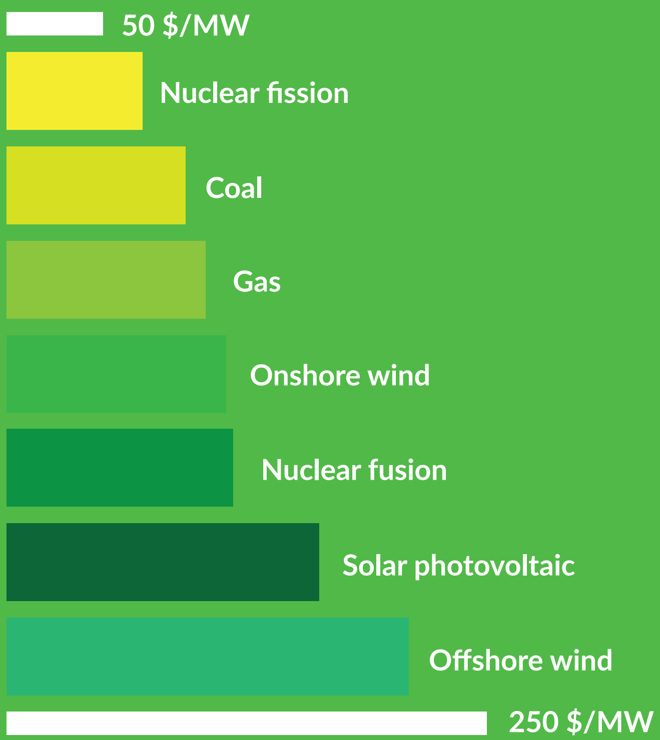

Figure 5.7. Bar graph comparing the Levelized Cost of Electricity of different fuel sources. Data: Entler, 2018.

the situation in Fukushima showed in 2011. Meanwhile, a nuclear fusion reactor is intrinsically protected against such a meltdown. Any disturbance in the fusion plasma causes it to extinguish. The plasma cools down immediately and all fusion reactions stop. Intrinsically, a fusion reactor cannot have a meltdown. Likewise, proliferation is not possible because nuclear weapons require a lot of fissile materials which are simply not present in a fusion reactor.

The third issue is that of nuclear waste. It is a fact that nuclear waste is produced while operating a tokamak. Over the years, the entire inside of the reactor becomes radioactive due to the fusion reactions. However, compared to the waste produced by fission reactors, the waste produced by nuclear fusion is much **shorter-lived**. Most of the nuclear waste produced in a fusion reactor loses its radioactivity within 50-100 years, which is much faster than the 1000s of years of some of the waste streams from a fission reactor.

Still, even short- or medium-lived nuclear waste is dangerous and needs to be handled carefully and securely. And while nuclear waste is nowhere near as much of an issue for fusion as for fission, it can definitely not be ignored either.

It is a societal question whether the production of radioactive waste will be a show-stopper for nuclear fusion (and fission). However, that discussion should ideally not be limited to radioactive waste alone. There are other growing concerns about e.g. **electronic waste (E-waste)** resulting from (amongst others) discarded solar panels.

Summarising

In this chapter, we saw some aspects that might determine whether nuclear fusion will be a success in the 2060s and beyond, or not. Much of it depends on the actual state of the electricity system by that time. What costs are acceptable? What land use? Which waste streams? The future will tell.

5.4

What about the money?

In previous chapters we read about the electricity market of the future and what nuclear fusion could bring to the table. We discussed some advantages of nuclear fusion over conventional renewables. However, we haven't talked about the costs of fusion very much yet. It turns out that building the first reactors in particular will be a costly endeavour. In this chapter, we will see why.

In the first half of this chapter we will discuss what factors are driving the bulk of the costs for the first generation of fusion reactors. The distribution of costs can be found in figure 5.8. Almost three quarters of the total costs are **capital costs**. These are the costs that we incur before we ever produce a single joule of energy. We will identify some key technologies that are the main drivers of these costs. Specifically, we will make a small case study of the magnet system of a fusion reactor.

Considerably smaller parts of the costs are caused by **repairments**, closely followed by **operation and maintenance** and, at the end of a reactor's lifetime, **decommisioning**. By the end of the chapter we briefly touch on the subject of the **cost of money**. Using this concept, we show why the huge share of capital cost in the total cost distribution is a serious problem for fusion reactors.

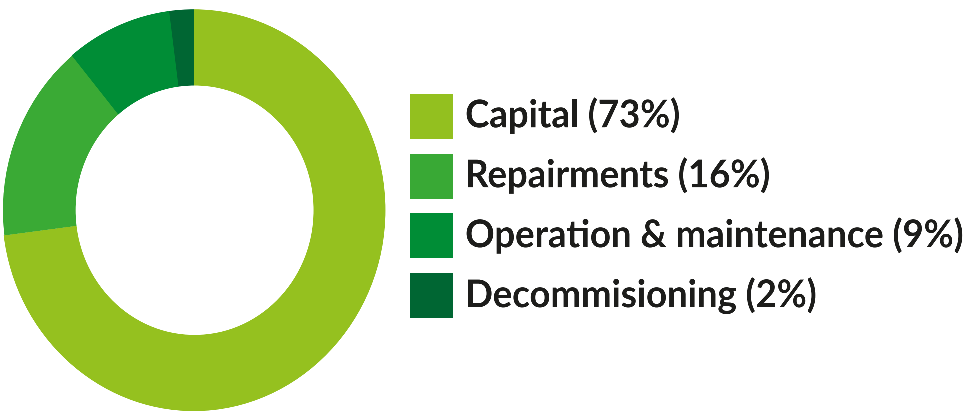

Figure 5.8. Pie chart of the distribution of costs in the categories capital, repairments, operation & mainetance and decommissioning. Data: Maisonnier, 2005.

Capital costs for the first generation

Three items on the list of capital costs stand out from amongst the crowd. These are the magnet system, the buildings and the heating systems.

Going in reverse order of importance, in third place we find that **constructing the heating systems** accounts for 8% of capital costs (or 6% of the total costs). While the tokamak is for sure the beating heart of a fusion reactor, it is only a blimp compared to everything else on the site. Take the ITER NBI (discussed in Module 3) as an example. From the perspective of the tokamak, this is just a hole in the wall that lets in heat and fuel. From the plant perspective, however, the ITER Neutral Beam Injector (NBI) facilities are much, much larger than the tokamak, as you can see in figure 5.9.

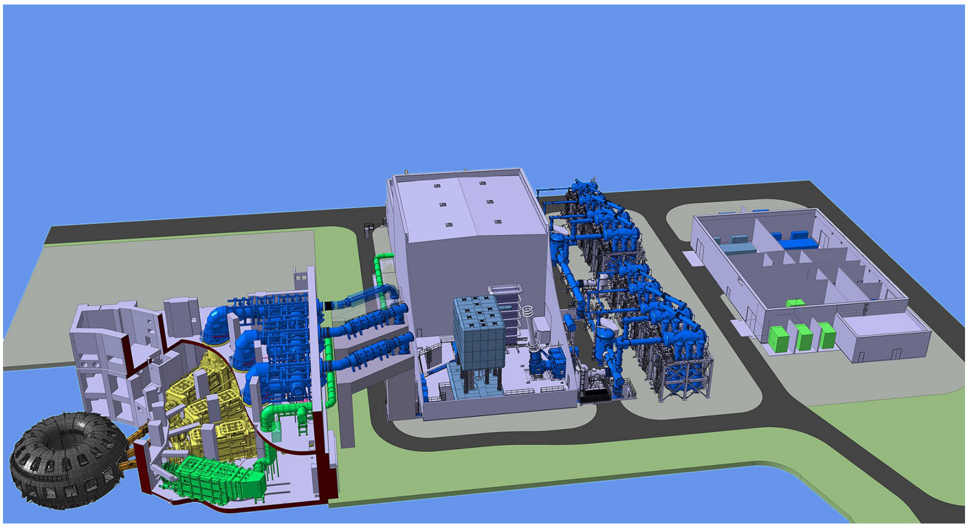

Figure 5.9. Schematic of the ITER NBI facility and how it connects to the vacuum vessel that contains the plasma. The size of the vacuum vessel pales in comparison to the entire NBI structure. Source: https://www.iter.org/newsline/-/3554

In second place, we find that the cost of **constructing all the buildings on site** amounts to 19% of the capital costs (or 14% of the total costs). To understand why, we only need to look at figure 5.10 which shows the entire ITER site from an aerial perspective. The buildings shown in figure 5.9 (including part of the tokamak building and the NBI facilities) are 'just' numbers (1) and (18). They are only a blimp on the map compared to all the other structures that are required to run a fusion plant. Building materials are expensive and so is the cost of labour. As a result, the cost of constructing building is huge. It is second only to one other thing: the magnet system.

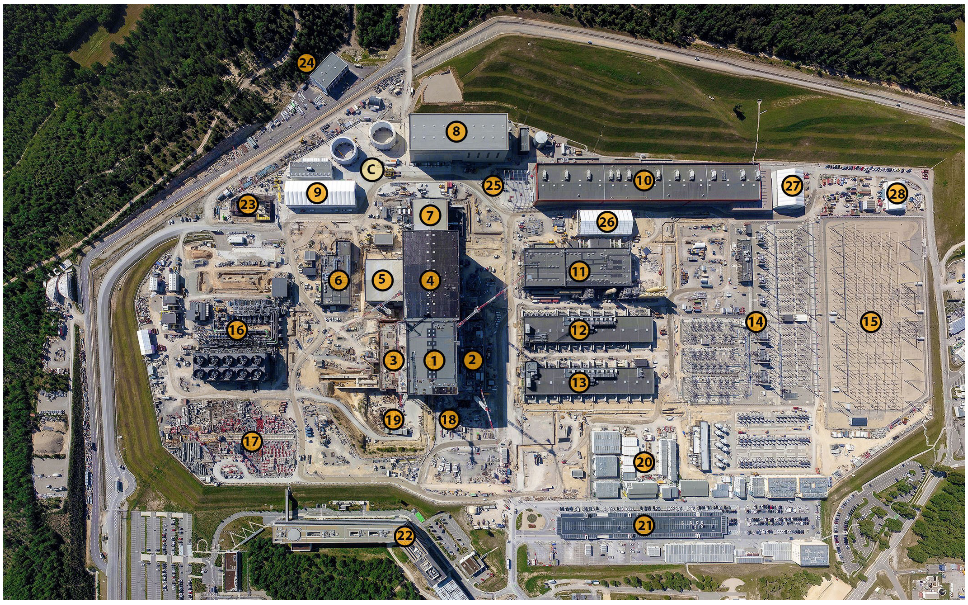

Figure 5.10. Aerial overview of the site with all the different buildings labelled:

(1) Tokamak Building
(2) Diagnostics Building
(3) Tritium Building
(4) Assembly Hall
(5) Radiofrequency Building
(6) Site Services Building
(7) Cleaning Facility
(8) Cryostat Workshop
(9) Magnets workshop
(10) Poloidal Field Coils Winding Facility
(11) Cryoplant
(12) Magnet Power Conversion Building #1
(13) Magnet Power Conversion Building #2
(14) 400 kV Electrical Switchyard (ITER Organization)
(15) 400 kV Electrical Switchyard (RTE, Réseau de transport d'électricité, France)
(16) Heat Removal Zone/Cooling Tower Zone
(17) Future location of the Control Buildings
(18) Future location of the neutral beam injection power supply
(19) Future location of the Hot Cell Facility
(20) Contractors Area
(21) Contractors elevated parking lot
(22) ITER Organization Headquarters
(23) Tokamak Assembly Preparatory Building (in construction)
(24) Assembly workshop
(25) Poloidal Field Coil Facility extension for cold tests
(26) Temporary storage
(27) Temporary storage
(28) Temporary storage
(C) Cryostat lower and upper cylinder (cocooned)

Credit: ITER Organization https://www.iter.org/newsline/-/3451).

The uncontested first place in the list of capital cost drivers, however, goes to the **construction of the magnet system**. At 42% of the capital costs, the magnet system is thus responsible for 31 cents out of every euro that is spent over the entire lifetime of a fusion reactor. Why is that the case? There are two main reasons. For one, they are amongst the largest and most powerful magnets ever to have been built. For another, they are superconducting. How do these two reasons - size and superconductivity - translate to capital costs?

Case study: why are magnets so costly?

At first glance, there is no problem per se with the magnets. The magnet industry already makes use of the superconducting materials that are being used for ITER and that will likely be used in future fusion power plants. These materials, like **NbTi (niobium-titanium)** and **Nb_3Sn (niobium-tin)**, are routinely used for the production of e.g. MRI scanners.

However, superconducting materials can be difficult to use in a production environment. While NbTi is easier to work with than Nb_3Sn, Nb_3Sn is favourable (and sometimes indispensable) from the operational perspective because it can generate larger magnetic fields and has a higher temperature threshold.

While such manufacturing problems have been previously tackled at the size of MRI magnets, they have not been tackled at the size of fusion magnets. As an example, there is the issue that Nb_3Sn is a brittle material. Brittleness is a material concept that we discussed in Module 4. Brittle materials break if you try to bend them (as opposed to ductile materials, which do in fact bend) and that makes them difficult to manipulate. Intricate processes are required to build such magnets and while technological solutions are routinely available at the size of MRI magnets, the same cannot be said for fusion-reactor-sized magnets.

The fact that fusion reactors require superconducting magnets brings a second challenge outside of the material one. Superconducting magnets have the strange (but very interesting) property that they become superconducting only below a certain temperature: the **critical temperature**. Above this temperature, they are almost perfect insulators which means that they conduct no electricity at all. The temperature threshold for Nb_3Sn to become superconducting is 18 K. For NbTn, the threshold is even lower at 9.5 K. This is very close to **absolute zero** at 0 K.

To use NbTn and Nb_3Sn, we thus need to cool them down to very low temperatures and, crucially, we need to maintain this temperature during the plant's operation. However, cooling near absolute zero (0 K) works very different from cooling near room temperature. Intricate cooling methods and very careful insulation are required to keep the magnets at this temperature. If we do not actively keep the magnets at this temperature, they would heat up to room temperature. To keep the magnets cooled, fusion reactors require a **cryoplant** (building 11 in figure 5.10), which (put crudely) is just a very big refridgerator.

This considerably adds to the expenses of the magnet system and thus to the capital costs, which helps put the magnet system squarely in first place as the most costly components of a fusion reactor.

Classroom Exercise 5.5
Kelvin (K) is a unit of temperature that is often used in physics and engineering instead of degrees Celsius (℃). The conversion between the two is straightforward. 0 K is equal to −273.15 ℃ and 273.15 K is 0 ℃. The formula is

$$T\,[\mathrm{K}] = T\,[°\mathrm{C}] + 273.15\ °\mathrm{C}$$

(a) What are the critical temperatures of NbTi and Nb_3Sn in degrees Celsius?

(b) How do these critical temperatures compare to the boiling points of liquid nitrogen (-196 ℃) and liquid helium (−268.95 ℃)? Which of the two do you expect to be cheaper as a coolant? And which do we need to cool NbTi and Nb_3Sn?

High-temperature superconductors (HTS) are considered to be a major breakthrough technology that could considerably accelerate the development of cost-effective fusion power plants. An example is rare-earth barium copper oxide (REBCO). It can generate much higher magnetic fields than NbTi and Nb_3Sn, which is favourable for the performance of a fusion reactor. However, that's not its only advantage. REBCO has a critical temperature of 92 K.

(c) Why is the critical temperature of REBCO an advantage?

Repairments

The second-largest contributor to the costs of fusion electricity is **repairments**. Although far smaller than the capital costs, repairments still contribute 16% to the total cost of a fusion reactor. In this category, we have the costs that come with repairing parts of the power plant.

Most of the repairments need to be made to parts of the wall and the divertor. We saw in Module 4 that these are continuously bombarded by neutrons and enormous heat loads. The walls of fusion reactors are lined by some of the most resilient materials ever to have been developed by mankind. Regardless, even these materials are spent after 2 to 5 years (depending on the specific component).

Replacing these parts is no easy task: the vacuum vessel in which the plasma is confined is radioactive, which means that we cannot simply send in people to replace the parts. Instead, robotic arms need to be used to replace the components. Due to the complexity of the parts and the reactor, this is a time-consuming task: it can take more than 6 months to fully replace all parts.

The remainder

Only 9% of the money is spent during the time that the fusion reactor is actually running. This category is called operation & maintenance and encompasses all the activities that are not undertaken during the building of the reactor or during major repairments. It includes all the personnel costs, fuel costs, electricity costs, minor maintenance, etc.

After operation & maintenance there is one final stage in the life of a nuclear fusion reactor, and that is its decommisioning. At some point the device becomes too old and needs to shut down permanently. This hails the start of the decommisioning phase, followed by the eventual dismantlement of the reactor. Even today, there is a lot of research being done in the nuclear sector on how best to decommision and dismantle nuclear devices. By the time that the first fusion devices will need to be dismantled, we will certainly have learned from the fleet of current-day fission reactors that will have since been decommissioned and dismantled.

Cost of money

In the introduction to this chapter we teased that the large capital costs for fusion could be a serious issue. This is due to the **cost of money**. Simply put, the cost of money is the **interest** we pay for the money we borrow. The interest we pay depends on the risk that investors take by lending us money and the time it takes for us to pay them back.

Investors need to pay the **73% capital costs up front** to build the reactor without knowing whether it will ever pay off, because there are no previous fusion success stories. They would be pioneering. And even if we could promise them 20% return on investment by the end of lifetime (fourty years from now), they would have to have to **wait for well over thirty years** to earn back their initial investment. If you were the investor, would you invest?

5.5 Can fusion become cheaper?

In the previous chapter, we discussed the main cost drivers for a fusion power plant. However, a discussion on how to reduce the costs cannot necessarily be followed along the same lines. Without expecting huge breakthroughs in technologies, like high-temperature superconductors or cheap alternatives to concrete, these costs will be what they are.

However, there are other ways in which we could think to reduce the cost of fusion electricity that do not require such monumental breakthroughs. In this chapter, we discuss three possible routes along which we might be able to optimise performance. All three routes are based on the same, simple thought: if the capital costs are what they are, we should work to **maximise electricity production**.

Why? Because if we can produce more electricity for the same initial investment (which is a large part of the total investment as we saw in chapter 5.4), the costs per MWh of fusion electricity will drop. Three knobs that we could potentially turn are **reactor availability (repairments and burn duty cycle)**, **plant efficiency** and **net electric output**.

By the end of this chapter, we will also touch on the concept of **technological learning**. It is not so much a knob that we can decide to turn ourselves, but rather an emerging effect. As we get more confident in building fusion reactors, building them will get cheaper and cheaper, particularly in the beginning.

Repairments

The reactor availability is the fraction of time that a reactor is actually operating. Say, we look at two reactors A and B. Reactor A is out of operation 5 weeks a year and thus has a reactor availability of around 90%. Reactor B, meanwhile, is out of operation for 3 months a year and thus has a reactor availability of 75%. All other things equal, reactor A produces 20% more electricity.

For fusion, reactor availability is a serious issue. One of the reasons for this follows directly from what we discussed in the previous chapters. **Repairments** of critical components in the vacuum vessel – like the first wall and the divertor that we discussed at length in chapter 4 – are very difficult and time-consuming. Replacing all the parts in the vacuum vessel can take as much as half a year. And if you need repairments every two to five years, at best, you're already spending 10% of your time on them, meaning your reactor availability drops to 90% at best. Solutions? **Longer-lived materials** and **faster remote maintenance** could seriously increase reactor availability and are therefore of prime importance if we want to decreases the cost of fusion electricity.

Burn duty cycle

A second effect that decreases reactor availability is the fact that current-day devices run in **pulsed mode**. That's not a choice, but a result of the reactor technology that is used. We saw in Module 1 that part of the magnetic field is generated inductively by a huge plasma current. This plasma current is, in turn, inductively driven by a large coil in the centre of the torus (central solenoid). The problem with this is that for induction to work, we require a **continuous change in magnetic field**.

While the reactor is operating, we are continuously increasing the magnetic field produced by the central solenoid. That's a problem. Not only do superconducting magnets have a critical temperature above which they cannot operate, they also have a **critical magnetic field**. At some point we cannot increase the magnetic field further and, as a consequence, cannot inductively drive the plasma current anymore. As this current tapers off, so does the confinement of the plasma come to an end.

Before the plasma can burn again, many systems need to be reset and that takes time. The fraction of time that the plasma is burning divided by the total time between consecutive burning plasmas is referred to as the **burn** duty cycle. Duty cycles are a general concept from engineering that quantify the fraction of time that a single system is operating during the entire on-and-off cycle. The duty cycle of one system (D) is found by dividing the runtime of that system (τ) by the runtime of the entire on-and-off cycle (T):

$$D = \frac{\tau}{T} \times 100\%$$

The burn duty cycle of the reactor strongly affects the total reactor availability. Assume for the moment that the burn duty cycle is also 90%. Then the total reactor availability, taking into account 10% repairments, drops to just over 80%. Increasing the burn duty cycle is thus another way to increase the availability and reduce costs.

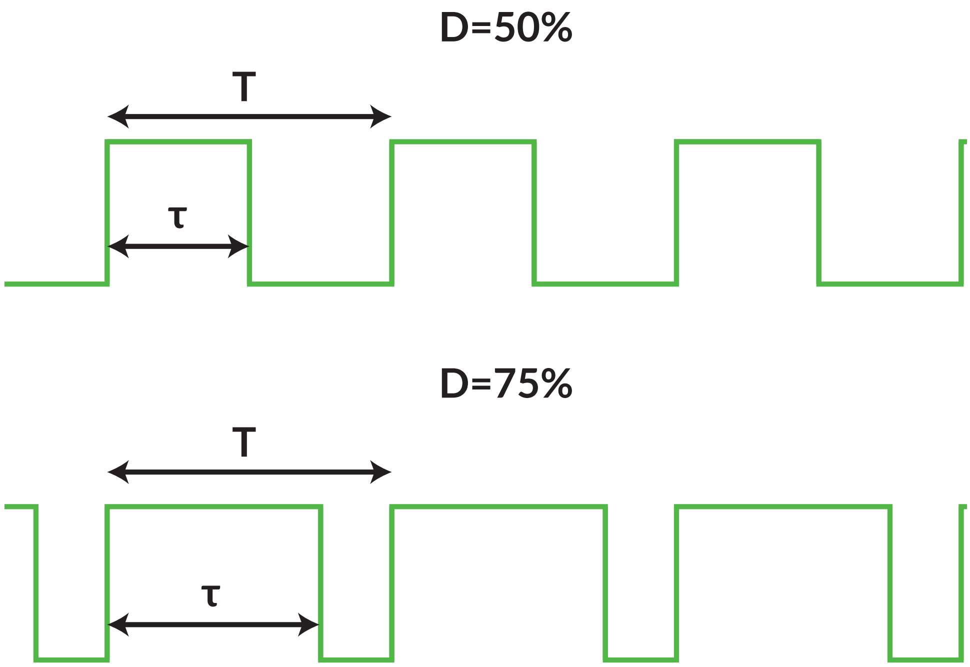

Figure 5.11. Two examples of duty cycles. In the top figure, the burn time τ is exactly half of the total cycle time T so $D = 50\%$. In the bottom figure, the burn time τ is three quarters of the total cycle time T so $D = 75\%$.

Plant efficiency and net electric output

A second factor that determines the cost of electricity is the plant efficiency, which is the ratio of the net electric power and the fusion power

$$\eta = P_{electric} / P_{fusion}.$$

In other words, it indicates how much of the energy released by the fusion reactions is actually converted into

Classroom Exercise 5.6

The **burn duty cycle** is important in determining the reactor availability. The standard scenario for ITER – called ELMy H-mode – will have a burn phase of around 300–500 s. The time between consecutive burns is around 1800 s.

(a) What is the maximum burn duty cycle of ITER in the ELMy H-mode scenario?

Between consecutive burns the ITER systems need to be reset. The plasma stops burning (but is still on for a time yet), while the plasma current is slowly reduced and the magnet systems are reset. At some point the plasma switches off only to be rekindled moments later. A new plasma current builds up and the auxiliary heating systems (see Module 3) heat the plasma back up to burning temperatures.

(b) What is the minimum duty cycle related to the reset (i.e. everything that is happening in between two burn phases) in the ELMy H-mode scenario?

The standard scenario for ITER is rather conservative. It is based on 'known' or at least expected behaviour of the fusion plasma. Part of the research in ITER will involve trying to increase this duty cycle. By extending the ELMy H-mode scenario, researchers are hoping to run the plasma for up to 2000 s at a time.

(c) What are the maximum burn duty cycle and reset duty cycle of ITER in the extended ELMy H-mode scenario? Hint: do you expect the reset time to change?

To get rid of issues related to pulsed operation, some researchers hope to one day have a steady-state fusion reactor that can keep the plasma burning indefinitely. However, with current-day technology such a steady-state reactor is unlikely to be cost effective. An alternative is to retain cyclic operation but with a very long burn phase, say 8 hours.

(d) What are the burn duty cycle and reset duty cycle for a steady-state fusion reactor? And for a cyclic reactor with an 8-hour burn time?

(e) What is the total reactor availability in both cases assuming that repairments account for 10% downtime?

electric energy. For the first generation of power plants, it is expected that $\eta < 20\%$, but in the longer run, the efficiency is expected to increase. However, to increase the efficiency, we need wall elements that can operate at higher temperatures and cooling materials that can sustain these temperatures as well and both of these still need to be developed. Increasing the plant efficiency is another way to increase the electricity output while not increasing the capital costs.

A third factor, that is similar to but not quite the same as plant efficiency, is the net electric output. The plant efficiency focuses solely on the amount of electricity that we can get out of the fusion plasma. However, we cannot sell all of this electricity. There are a lot of systems that consume electricity inside a fusion plant, for example for plasma heating, magnet cooling and coolant pumping. All the power that is used to keep these machines running is subtracted from the electric power generated by the plant before anything can be delivered to the grid, so

$$P_{net} = P_{\text{electric}} - P_{\text{plant}} = \eta \times P_{\text{fusion}} - P_{\text{plant}}.$$

For the first generation of fusion power plants, this consumption is expected to be about 50% of the gross electricity production. This means that only 10% of the fusion power is eventually delivered to the grid. Increasing this fraction is a surefire way to decrease the cost of fusion electricity.

Technological learning

One final aspect related to the cost of electricity is the concept of technological learning. It is less of a knob that we can actively turn and more of an emerging property. As people start to find their footing in the field of building and operation fusion reactors, they will develop ways of doing things more cleverly, more cheaply and more efficiently. One of the main cost drivers of the first reactors, as we saw, is the magnet system. Much of the costliness of the magnet system comes from the fact that they are a first-of-a-kind technology. However, as we build more and more of these magnets we become better equipped to build them.

Concepts like these are formalised in the study of technological learning. The speed at which we build up experience can be expressed mathematically by the so-called **progress ratio (*PR*)**. The progress ratio describes the reduction in cost after each **doubling** of the total production. If the PR of building a fusion power plant is

0.9, the second power plant will be ten percent cheaper than the first, the fourth power plant ten percent cheaper than the second and the eight power plant again ten percent cheaper.

The total reduction in cost is called **learning factor (LF)**, which tells us that the cost of a power plant after N doublings (i.e. the cost of power plant number 2^N) is

$$LF = PR^N.$$

As we build more reactors over time, N increases and so does LF. However, the number of reactors that we need to build to double the entire reactor fleet also increases over time. To go from one reactor to two reactors, we need to build one extra. To go from two reactors to four reactors, we need to build two. To go from four reactors to eight reactors, etc. So the speed of learning decreases over time. Still, the first few generations of reactors will quickly become cheaper to build thanks to this effect.

In summary

Long story short, by generating more electricity with the same fusion reactor, we can make electricity cheaper. An important role is played by the reactor availability, which is dominated by downtime due to repairments and by the burn duty cycle. Two other factors of importance for the cost of fusion electricity are the plant efficiency and net electric output. They determine how much of the produced fusion power actually ends up in the electricity sockets of our homes.

Finally, we saw that technological learning will likely make fusion reactors cheaper to build as time goes by. So if we can just get started on that first reactor, each one that we build afterwards will be easier to build.

5.6 ITER and onwards

Nuclear fusion has never lacked issues and challenges. In the very beginning, those challenges were often related to the (plasma) physics of fusion. However, as the decades passed and many of the bits and pieces of physics were sorted out, challenges of a new type arose. As devices got bigger and magnets stronger, the materials they were built out of were strained to (and sometimes strained beyond) their maximum. Nowadays, it also has some of the most complex issues in technology, engineering and manufacturing. Fusion has long since stopped being 'just' a physics experiment.

In this module we marched ahead of the troops for a bit and it was a very interesting exercise. Sometime in the coming decades, nuclear fusion will not only have stopped being a physics experiment – **it will cease to be an experiment** entirely. Rather than being the end of issues and challenges, this means that we get a whole new class of trouble to centend with: that of an economical nature.

That does not mean that to work in fusion, you need to be an economist though. Far from it. While many of the new challenges will be driven by economical motives, their solutions are still largely to be found in the realms of physics and engineering. Besides, if you're now in highschool, we only need to look five to ten years in the future to see what opportunities will be there. In this last chapter of this last module, we will go over the three up-and-coming European devices: **ITER**, **IFMIF-DONES** and **DEMO**.

ITER

We've mentioned **ITER** (Latin for 'the way') on quite a few occasions throughout this and other modules. Notably, in module 2 we learned a lot about its history.

Now, what's in the near future? According to the current planning of ITER (date of writing: 14 september 2022), they're still on their way to finishing the build of the reactor by December 2025, which is when they want to turn the device on for the first time to produce the **first plasma**.

However, the first fusion plasma (deuterium-tritium) is not expected until the mid 2030s. Before they go nuclear, they'll want to be super sure that everything is working as planned. After all, once they go nuclear, there's no going back into the (now radioactive) vacuum vessel. There are many research lessons to be learned in ITER. Two that stand out are to **control plasma instabilities** and to **breed tritium**. We discuss tritium breeding as the last aside of this module. So, there is enough work to be done in ITER provided that you graduate halfway through the next decade.

IFMIF-DONES

Another huge fusion experiment is the **International Fusion Materials Irradiation Facility – DEMO Oriented Neutron Source** (IFMIF-DONES). The **IFMIF** part of the name tells us that this facility will be used to study the effects of irradiation on materials that we will want to use in a fusion reactor. This mainly involves studying the effects that neutrons have on the plasma-facing components: they cause a lot damage and render materials radioactive. Neither is a very pleasant effect (see module 4).

Nowadays, studies of this type are still carried out in experimental fission reactors. However, those are not ideally suited to generating the enormous number of high-energy neutrons produced by ITER's successor, DEMO (see next section). We also cannot use ITER for this kind of study, because the neutron flux in ITER is roughly a factor of 100 smaller than in DEMO.

A component with a 2-year lifetime in DEMO, would have to be irradiated for 200 years in ITER to see similar levels of radioactivity and damage. This is where the **DONES** part comes in. IFMIF-DONES will be simulating the kind of neutron flux that plasma-facing components are subjected to in a full-fledged fusion power plant.

ASIDE

Tritium breeding

Tritium breeding is definitely amongst the next big steps in nuclear fusion. Whenever in fusion we say that there's enough fuel for the next thousand years, there's a big if. There's enough fuel if and only if we get tritium breeding to work. As it stands, there is far from enough fuel in the world. And because tritium (**T** or **^{3}H**) is unstable (i.e. radioactive), each 12.3 years half of it decays following

$$^3\mathrm{H} \rightarrow {}^3\mathrm{He} + e^- + \bar{\nu}_e + \text{energy}.$$

One tritium (or 3-hydrogen) can decay into one 3-helium while also emitting an electron **e^-** and an electron antineutrino **$\bar{\nu}_e$** along with some energy. The fact that energy is emitted as a result of this decay (i.e. that the decay is exothermal) means that it can spontaneously happen. As a result, no tritium is available in nature. Moreover, even the artificial production of tritium is almost non-existent. Currently, there is about 20 kg of tritium worldwide.

Long story short, nuclear fusion reactors will need to generate their own tritium. And the early reactors will need to produce more, because they will use up almost the entire inventory and will thus also need to supply the tritium for future generations of reactors. Making (or breeding) tritium is crucial for fusion.

Lithium

Research into the breeding blanket has been ongoing for quite a while now. Several alternative concepts have been proposed and are being worked out as we speak. However, the proof of the pudding is in the eating. We'll only be able to test these blanket modules once ITER is up and running. All of these technologies, however, are based on the same material: lithium.

The reason for that is quite simple. If a neutron reacts with lithium, one of the products is tritium. We get the neutrons for free in the deuterium-tritium (DT) reaction

$$^2\mathrm{H} + {}^3\mathrm{H} \rightarrow {}^4\mathrm{He} + n + \text{energy}$$

so we only need lithium to have the following reactions

$$^{6}Li + n \rightarrow {}^{3}H + {}^{4}He + energy$$

$$^{7}Li + n \rightarrow {}^{3}H + {}^{4}He + n - energy$$

The 6-lithium reaction is clearly favourable, because it is exothermal. Not only does it yield a tritium, it also produces some energy that we get for free on top of the energy produced by the DT reaction. Unfortunately, 92.5% of all lithium in nature is 7-lithium, which gives us an endothermal reaction (i.e. it costs energy).

Additionally, each DT reaction requires one T and yields one n. Each n can be used to produce exactly one new T, but only if we capture all of the neutrons. We can't, because in some places of the reactor there is no room for a blanket, so there are holes through which some neutrons escape. And even if we could capture all the neutrons, that still leaves us with no excess T for new power plants.

Neutron multiplication

For this reason, we also want to multiply the neutrons. We need a material that, if given a neutron, gives us two neutrons back. Two materials are suitable for our purposes, beryllium (Be) and lead (Pb). They react as

$$^{9}Be + n \rightarrow 2\ {}^{4}He + 2\ n - energy$$

$$^{208}Pb + n \rightarrow {}^{207}Pb + 2\ n - energy$$

We can see that both of these are endothermal, so the reactions need energy to take place. For lead, the needed energy is three times as high as for beryllium. But we can make sure this energy is provided: if we place the material close enough to the plasma (so not too deep in the wall), the neutrons will still be quite fast, meaning they have enough energy to start the reaction. By making use of a clever design with the neutron multipliers, we expect that we can produce roughly **15% more tritium** than is consumed. In ITER, we will investigate different structures for tritium breeding, including both neutron multipliers, to see what works best.

ASIDE

Excess tritium is needed

We've already mentioned that excess neutrons are required to offset losses in parts of the reactor that don't have a blanket. We also know that we need excess tritium for future generations of reactors. But we don't need it for that alone. Summarising, three main issues are

(I) the natural decay of tritium, that causes the loss of more than 5% of our tritium fuel every year;
(II) the retention of tritium within all the piping in the reactor and to the tritium treament plant;
(III) a start-up inventory for new reactors, as mentioned before.

These three main issues mean that the extra 15% of tritium that we can breed may just be enough to actually achieve fusion energy. This makes it one of the most crucial aspects of research to be conducted at ITER.

Moreover, tritium breeding is crucial in the context of deployment. While the main focus in this module was on the costs of fusion, another aspect of deployment is the speed at which we can build new power plants. While we won't dive into that particular topic in this module, you could imagine that the rate at which we are able to produce excess tritium for new start-up inventories limits the rate at which we can build new reactors.

DEMO

We have come across this name several times now: the **DEMOnstration power plant** is going to be ITER's successor. As such, it will use the lessons we have learnt in the previous devices. The main goal of DEMO is not, however, experimental. Rather, DEMO is going to demonstrate the production of electricity with nuclear fusion. The goal is to deliver a net power of

$$P_{\mathrm{DEMO}}=300\text{–}500\ \mathrm{MW}$$

to the grid. This does not mean that no research will be conducted at DEMO. It will in fact be used to optimize the design of a fusion reactor, but primarily from an operational standpoint. How do the material choices hold up in a full-fledged reactor? How easy is it to replace them? Is remote maintenance difficult? Enough to do in DEMO as well!

Right now, people are already working very hard towards conceptualising what DEMO should look like. Several options are on the table and it is not even clear whether DEMO will be a tokamak like ITER. Much progress is being made in Wendelstein 7-X, which is a European stellarator experiment that has been in operation for several years already and shows promising results. Sadly, we **cannot wait for ITER's deuterium-tritium** results before deciding what our DEMO should look like. It takes so much time to develop the plans for a fusion power plant and to eventually build it, that we need to make many choices right now. Luckily, there's decades of experience in smaller and larger fusion experiments worldwide.

Other DEMOs

These three devices are major players on the EUROfusion roadmap to Fusion Energy (see further reading). However, there's fusion industry beyond Europe as well. ITER is a collaboration between seven partners worldwide. Our DEMO, on the other hand, will be developed in a partnership between the EU and Japan. China, India, South-Korea and the US have plans for their very own DEMO reactors, while Russia is trying to develop a fission-fusion hybrid.

Further reading

The European fusion plans are described in detail in the **EUROfusion Roadmap to Realising Fusion Energy**. While it will definitely contain things that were not discussed in our five modules, you will be surprised how much of it you will understand after having worked through all five. And on the somewhat nearer term, the **ITER milestones** are also interesting to look at!

https://www.euro-fusion.org/eurofusion/roadmap/

https://www.iter.org/proj/itermilestones

Colophon

Deployment is the fifth of five modules on nuclear fusion. They teach about basic concepts in nuclear fusion at the secondary school level. All modules can be found free of charge on the website of FuseNet, **https://fusenet.eu**.

Authors
Bas Marechal, Sander Korteweg and Jens Peter Frankemölle

Editors
Sander Korteweg and Sjoukje Tijmensen-Hoekstra

Graphical design
Jens Peter Frankemölle

Cover image
Bruno/Germany via pixabay.com

Publisher
FuseNet

Financial support
This project would not have been possible without the financial support of EUROfusion.