Nuclear rockets 8

In a chemical rocket, the kinetic energy of the propellant is obtained from chemical bonds by burning the fuel. The available energy per unit mass is limited to about $10\,\mathrm{MJ\,kg^{-1}}$, which limits the effective exhaust velocity to about $4.5\,\mathrm{km\,s^{-1}}$ (equation 3.24). Nuclear reactions, on the other hand, can release the much larger binding energy between nucleons, of the order of $10^8\,\mathrm{MJ\,kg^{-1}}$, implying effective exhaust velocities of around $0.05\,c$. In this chapter we look at the basics of nuclear fission and fusion, including reactors and the optimal reactions for propulsion. We then examine the various ways in which these can be used to propel a rocket, which includes both sustained reactions and miniature explosions. Although nuclear-propelled rockets have not been used in space (as far as we know), there has been quite some research and development. We will also look at the principles of antimatter propulsion and see what effective exhaust velocity this could achieve.

8.1 Nuclear binding energy

Nucleons – protons and neutrons – are held together in the nucleus of an atom by the strong nuclear force. The attractive nature of this force means that as individual nucleons are brought closer together, the force does work, and this energy – the binding energy – is released. As more nucleons are added, not only does the total binding energy increase, but so does the binding energy per nucleon, because each new nucleon interacts with all of the existing ones. However, in addition to this attractive nuclear force there exists also the electromagnetic force that repels the positively charged protons, and this works to lower the binding energy. For small nuclei the nuclear force is the dominant one. But as the nuclear force has only a very short range, the electromagnetic force lowers the binding energy per nucleon for larger nuclei.[1]

The exact physics is more complicated than this, but the overall result is that as the number of assembled nucleons increases, the binding energy per nucleon at first increases, and then decreases, as shown in figure 8.1. The combination of two small nuclei therefore leads to a net release of binding energy, because the resulting product will have more binding energy than the sum of the original components. This is the process of nuclear fusion. The maximum binding energy per nucleon occurs at nucleon number $A = 62$, corresponding to nickel-62, with 8.79 MeV per nucleon. Fusion of nuclei larger than this no longer releases energy, because the resulting product would have less binding energy than the sum of the original components. Instead, breaking up a single larger nucleus into two or more products releases the excess binding energy. This is the process of nuclear fission.

[1] At very short distances the nuclear force is repulsive, which is what prevents the nucleus from collapsing.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

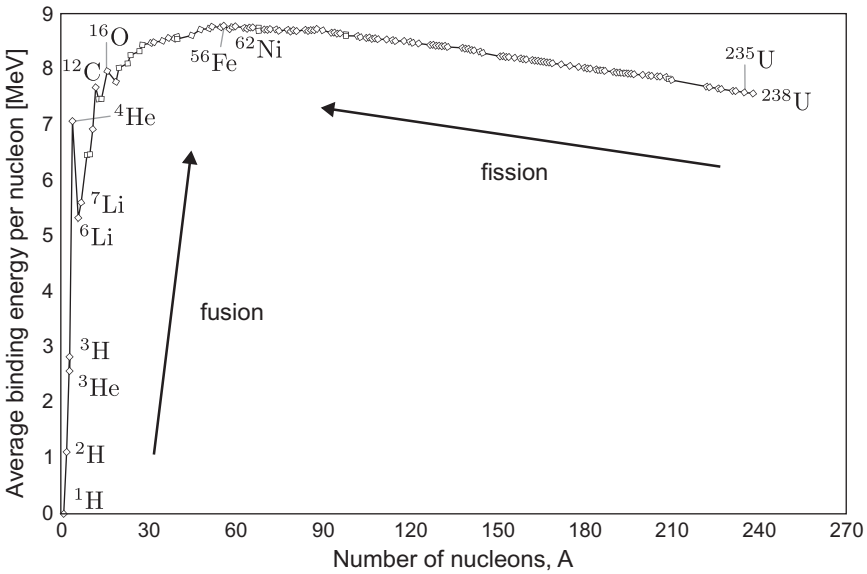


Fig. 8.1 The variation of the binding energy per nucleon in atomic nuclei with the number of nucleons (A, also known as the atomic mass number) for a selection of nuclei. Assembling small nuclei releases energy (fusion), as does breaking up large nuclei (fission). Credit: Public domain, from Wikimedia Commons, modified by the author.

Figure 8.1 shows that the typical binding energies per nucleon are of the order of several MeV per nucleon; 1 MeV per nucleon is equivalent to about $0.96 \times 10^8\,\mathrm{MJ\,kg^{-1}}$. For comparison, the rest energy ($E_0 = mc^2$) of a proton is 938.27 MeV and that of a neutron is 939.57 MeV.

8.2 Nuclear fission

Nuclear fission is the process by which a heavy nucleus splits into two or more lighter nuclei. It can occur spontaneously, but the probability of it occurring is often larger if a neutron with a certain range of energies interacts with the nucleus. This is relevant because fission usually releases unbound neutrons in addition to nuclei. When, on average, more than one of the released neutrons goes on to cause another fission event, there will be an exponential increase in the number of fission events for as long as there is still fissionable material available.

The key to extracting energy from a fission reaction is therefore to control the number of neutrons. The *neutron multiplication factor* k is the ratio of the number of neutrons released in a generation of reactions to the number in the previous generation. When $k > 1$ we get an exponential increase; the reaction is said to be supercritical. As the timescale between fission events is very short, this leads to a very rapid release of energy. This is the goal of a fission bomb.

It would seem that in order to control the rate of energy release, we would need to keep k very close to one for an extended period of time. Fortunately for the peaceful extraction of energy, neutrons are released not only during the rapid fission event itself. The product

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

Table 8.1 Example uranium-235 fission products showing the amount of energy released that ends up in the kinetic energy of these products as well as neutrons and photons. Data from De Sanctis et al. (2016). Decays to three or more product nuclei are also possible, but are much less probable.

Product	Product	No. neutrons released	Energy released [MeV]
$^{144}_{54}$Xe	$^{90}_{38}$Sr	2	176
$^{141}_{56}$Ba	$^{92}_{36}$Kr	3	173
$^{141}_{55}$Cs	$^{93}_{37}$Rb	2	180
$^{137}_{53}$I	$^{96}_{39}$Y	3	179
$^{136}_{52}$Te	$^{97}_{40}$Zr	3	182
$^{127}_{50}$Sn	$^{105}_{42}$Mo	4	178

nuclei resulting from the fission are generally unstable and undergo radioactive beta decay. The resulting nucleus is sometimes excited and may then emit a neutron. Some of these neutrons have the right energy to cause fission, and this contributes to the neutron budget and thus the value of k. The fortunate aspect is that because beta decay is controlled by the weak interaction, its timescale is orders of magnitudes longer than the timescale on which the so-called 'prompt' neutrons are released directly from the fission. The 'delayed' neutrons following beta decay are emitted of the order of 0.01–100 s after the fission. This is long enough to implement a control mechanism to react to the number of neutrons available and thus keep the reaction near critical ($k = 1$). This is achieved in a fission reactor with control rods, which are made from strong neutron absorbers such as boron or cadmium. By moving these in and out of the volume in which the fission is taking place, the reaction rate can be controlled. It is the existence of the delayed neutrons that makes it possible to control nuclear fission.

The most widely used element for harnessing energy from fission is uranium-235, which contains 92 protons and 143 neutrons. When a neutron hits the nucleus, it can be absorbed to create uranium-236. This is unstable and splits into two (occasionally more) product nuclei. There are many possible paths for this decay, as shown in table 8.1, and these occur randomly in different proportions. The product nuclei tend to have unequal masses, the heavier one with a nucleon number A between 125 and 150, and lighter one with A between 80 and 105. Each fission typically releases two, three, or four neutrons, plus around 200 MeV of energy. This energy is the difference between the binding energy of the parent nucleus and that of the products. The parent has the order of 7.6 MeV per nucleon, whereas the products have the order of 8.5 MeV per nucleon, a difference of about 0.9 MeV per nucleon with 235 nucleons. This binding energy is released in the form of both high-energy photons (gamma rays), which receive about 7 MeV, and kinetic energy of the product nuclei (165 MeV) and neutrons (5 MeV). The rest of the energy comes from the beta and gamma decays of intermediate fission products in the form of photons (6 MeV), electrons (7 MeV), and antineutrinos (10 MeV). This released energy of 200 MeV compares to a rest energy of

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

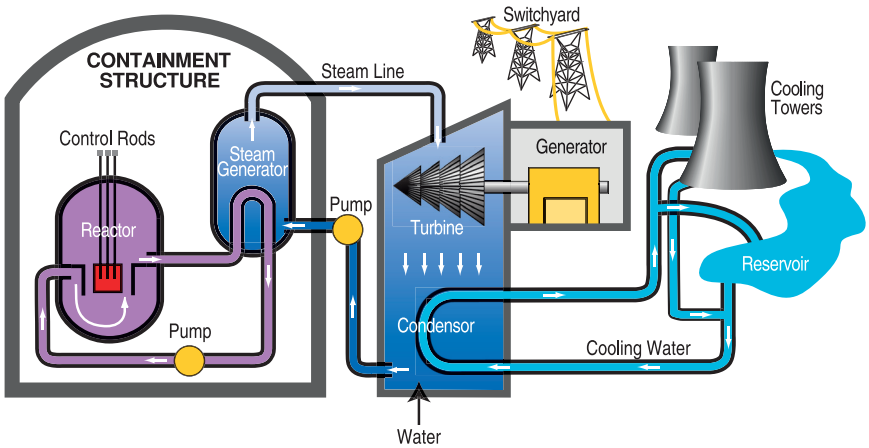


Fig. 8.2 A pressurized water nuclear fission reactor. Credit: Tennessee Valley Authority, public domain, from Wikimedia Commons.

the original uranium-235 nucleus of 2×10^5 MeV. Thus about 0.1% of the initial mass–energy is converted into kinetic energy of the fission products.

In a nuclear reactor, the energy of the gamma rays and the kinetic energy of the product nuclei are absorbed by a working fluid (often water). This drives a heat engine, usually a steam turbine, to produce electricity in a closed cycle. The typical layout of such a nuclear power station is shown in figure 8.2.

The neutrons released by a uranium-235 fission typically have kinetic energies of 2 MeV each. At this energy they have a high probability of causing fission in uranium-235, so most could contribute to the value of k. However, naturally occurring uranium on the Earth is a mixture of isotopes, with 99.3% uranium-238 and only 0.7% uranium-235 (plus 0.006% uranium-234). Not only does uranium-238 have a very low probability of undergoing induced fission, it also strongly absorbs neutrons with energies around 2 MeV. Thus most of the neutrons released by a uranium-235 fission would be absorbed by the uranium-238 isotopes, and very few would go on to cause further fission in uranium-235, resulting in a subcritical reaction. For this reason a sustainable nuclear fission reaction cannot spontaneously occur in natural uranium.

To achieve sustainable fission we have two options. First, we could increase the proportion of uranium-235 – by removing uranium-238 – so that enough neutrons remain unabsorbed and so attain $k > 1$. In principle – in an infinite mass of uranium such that neutron losses from the surface can be ignored – this can be achieved once the uranium is enriched to 5.5% uranium-235. In practice, though, an enrichment of 20% or more is used in what is then called a 'fast reactor', because it uses the high-energy neutrons directly.

A second option for achieving sustainable fission is to exploit the fact that the probability of a neutron causing fission, as opposed to being captured, depends on the neutron's energy. Neutrons with an energy significantly lower than 2 MeV are more likely to cause fission in uranium-235 than be captured by uranium-238. We can therefore use a *moderator* to slow down the neutrons, via elastic collisions with other atoms or molecules, to energies of around 0.1 eV (a temperature of 800 K) or less. Graphite (carbon) and heavy water (where

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

deuterium, ^{2}H, takes the place of ^{1}H) are good moderators, and some nuclear reactors use these to achieve fission with natural uranium. Normal water ('light water') is a good neutron moderator, but unfortunately it is also a good neutron absorber, so we cannot achieve a stable fission reaction with natural uranium and normal water. However, if we enrich natural uranium to have more than 3% uranium-235, this is sufficient to produce enough neutrons to compensate for the absorption.[2]

Most commercial reactors today use low-enriched (3–5%) uranium as the fuel and light water as the moderator. Uranium enriched to about 20% is used in the reactors of nuclear submarines, because such reactors are more compact and allow the submarine to operate for longer without refuelling. Fission bombs use uranium enriched to 80% or more, to provide a much larger k value and thus a very rapid release of energy.

Uranium is not the only fissile material available. Plutonium (atomic number $Z = 94$), specifically the isotope plutonium-239, has been used for both reactors and bombs. It is barely naturally occurring (half-life of 24 kyr), but is produced in uranium reactors when uranium-238 captures a neutron. In fact, as the fuel ages in a normal uranium reactor, plutonium-239 becomes responsible for a significant fraction of the energy production.

Controlling the neutron flux is central to controlling a fission reaction, and here the shape of the fissile material plays an important role. If the material is arranged in a long thin rod, it has a large surface area for its volume, so a lot of neutrons escape from the surface before they can cause a fission event. If we instead arrange the fuel into a shape that minimizes the surface area for a given volume, which is a sphere, then a larger proportion of the neutrons are captured. The surface area to volume ratio of a sphere of radius r is $1/r$, and the rate of neutron loss is proportional to this. Thus for material that in principle could achieve criticality ($k > 1$), a sufficiently small sphere will actually be non-critical. The *critical mass* is the minimum mass that a sphere of fissile material (of given density) must have in order to achieve criticality. For pure uncompressed uranium-235, the critical mass is about 50 kg, corresponding to $r = 8.5$ cm. This means that two isolated spheres each slightly below this mass are both subcritical. But if we bring them close together, they will turn critical and explode uncontrollably. This is in fact the basic approach to detonating a fission bomb (see section 8.3.3).

Even without the goal of generating an explosion, the loss of neutrons from the surface of fissile material means that the most compact reactors should aim to minimize their surface area for a given volume. A spherical reactor is not very practical, however, so they tend to be cylindrical instead. A light water reactor produces energy at a rate of about 30 MW per tonne of fuel, with one tonne generating around 40 gigawatt days of energy. This corresponds to 3.5×10^6 MJ kg^{-1}. This is much less than the 80×10^6 MJ kg^{-1} (200 MeV/235 m_u) released by the fission reaction itself, because most of the fuel is uranium-238 that is not consumed. A typical reactor contains about 100 tonnes of enriched uranium fuel and produces around 3 GW of thermal power, of which around 1 GW can be converted into electricity by a steam turbine.

[2] Around 1.7 billion years ago, uranium-235 – which has a half-life of 704 Myr – made up about 3.1% of naturally occurring uranium, and so with a normal water moderator fission would have been possible. This actually occurred at that time in what was once a uranium-rich deposit in Oklo, Gabon, after groundwater seeped in to act as a moderator. This was identified after geological samples showed a significantly lower level of uranium-235 than the global average and other nearby mines, indicating fission had already occurred.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

8.3 Fission rockets

We saw in the previous section that the fission of one uranium-235 nucleus puts about 165 MeV of kinetic energy (K) into the product nuclei. If all of these nuclei (of combined mass $m \simeq 233 m_u$) could be directly used as rocket propellant, they would provide an effective exhaust velocity of $\sqrt{2K/m} = 11\,700\,\mathrm{km\,s^{-1}}$, or $0.039\,c$. How this translates into the Δv of a rocket depends on how this energy is actually harnessed. Here we look at different ways to do this.

8.3.1 Thermal fission rockets

Perhaps the simplest use of the fission energy is to use it to heat an inert propellant, as shown in figure 8.3. The fission occurs continuously in a controlled manner in a reactor similar to a commercial reactor for generating electricity. But instead of the reactor heating a working fluid in a closed cycle to drive a turbine, the energy heats a propellant that is ejected. This delivers a much lower effective exhaust velocity than using the fission products directly, but a much larger thrust. The difference with respect to a chemical rocket is the source of the heat. After that, the rocket behaves just like any other thermal rocket discussed in chapter 5, in particular by using a nozzle to maximize the conversion of heat into kinetic energy.

A typical propellant is hydrogen on account of its low molar mass, thus giving a large effective exhaust velocity (equation 5.9). The hydrogen does not burn, and the exhaust is not radioactive provided all of the fission decay products remain within the fuel rods. In practice, however, corrosion of the fuel rods puts radioactive contaminants in the exhaust.

To keep the mass of the rocket as low as possible, the engine should use highly enriched uranium as fuel to avoid carrying non-fissile uranium-238. This has the further advantage that we do not need a moderator to sustain a reaction (a fast reactor), giving a further mass saving.

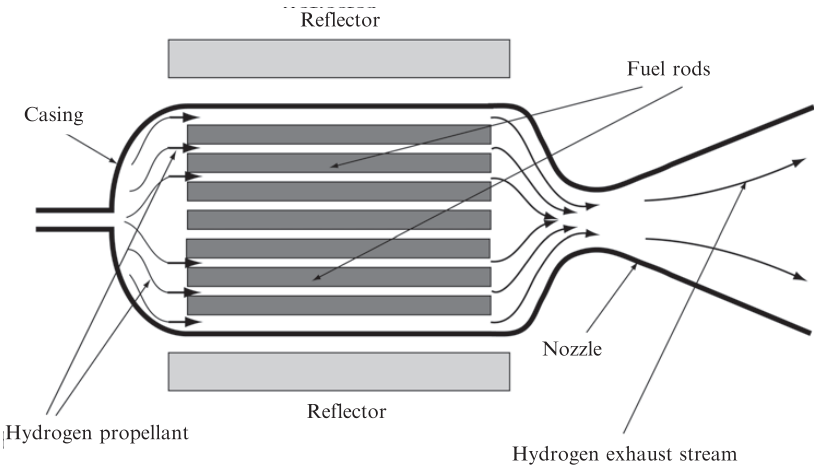


Fig. 8.3 A thermal fission rocket, in which a fission reactor heats an inert (non-combusted) propellant. Credit: Used with permission of Springer Nature BV, from Turner (2009), conveyed through Copyright Clearance Center, Inc.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009



A NERVA nuclear thermal rocket engine in testing. Credit: Photo courtesy of National Nuclear Security Administration / Nevada Site Office, public domain, via Wikimedia Commons.

Fig. 8.4

Although this makes the reactor smaller, it also increases the neutron loss as discussed above, lowering the reaction rate. For this reason, neutron reflectors are often employed in small reactors to compensate. A neutron reflector is a material that scatters neutrons elastically – but not specularly – so it needs to have a low absorption cross-section and not be fissionable. Various materials fulfil these requirements. Graphite and beryllium are frequently used in light water reactors, and have the advantage of low density. But they are also moderators that lower the energies of some of the neutrons, which is undesirable in a fast reactor. High-atomic-number materials are also good reflectors, but their high density makes them less desirable in a rocket.

To maximize the effective exhaust velocity we should try to achieve the largest reactor temperature possible (T_c in equation 5.9). This is often limited in practice by the melting temperature of the materials used to make the engines, and in the case of a fission reactor also the melting temperature of the fuel. Pure uranium melts at around 1400 K, which is rather low, so the fuel is usually made into a compound of uranium dioxide, which melts at 3075 K.

Thermal nuclear engines have been developed by both the USA and the Soviet Union/Russia. Between 1955 and 1972 the USA ran a programme called NERVA, Nuclear Engine for Rocket Vehicle Applications (figure 8.4). This developed working engines that were tested on the ground, although they never flew in space. Various versions were developed. One of them used a 1.5 GW reactor to heat a hydrogen propellant to $T_c = 2500\,\mathrm{K}$ and achieved $v_e = 8.7\,\mathrm{km\,s^{-1}}$. This velocity is only about twice that of a chemical engine, and far below the maximum achievable when directly using the fission products. The temperature is actually lower than that often achieved in a chemical thermal rocket engine. But because the propellant is much lighter – H_2 with a molar mass of two, instead of H_2O with a

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

molar mass of 18 – the effective exhaust velocity is nonetheless higher; equation 5.10 gives $v_e = 8.5\,\mathrm{km\,s^{-1}}$.

From the above velocity and engine power we compute the mass flow rate (equation 3.22) and thrust (equation 3.10) to be $40\,\mathrm{kg\,s^{-1}}$ and 345 kN, respectively, assuming all of the power is used. The engine itself had a mass of about 8000 kg; another 6000 kg was required for shielding the radiation. Although it used highly enriched uranium (over 20%), the reactor nonetheless used a moderator made from graphite and zirconium hydride. The engine was designed to operate for a few hours.

This approach to nuclear propulsion has a number of challenges. One is that even if we stop the reaction to turn off the engine, the fission decay products still produce heat. To avoid the core melting, we would have to pass hydrogen over it to cool it, which would still generate thrust. This makes it difficult to operate as a restartable engine. The decay product xenon, which is a strong neutron absorber, can also prevent the reaction being restarted (neutron poisoning).

The temperature of the nuclear core, and therefore the velocity of the propellant, is limited in the above designs by the use of solid fuel. Using a melted, liquid fuel permits higher temperatures. One approach is to confine the fuel within rotating tubes that prevent the hottest parts of the fuel touching the surface of the tubes. In this way a hydrogen propellant temperature of 5500 K could be achievable, providing $v_e = 17\,\mathrm{km\,s^{-1}}$. Even higher temperatures are in principle achievable with a gas core reactor, in which the fuel is vaporized into a plasma that heats the propellant by radiation. Exhaust velocities of 30–70 $\mathrm{km\,s^{-1}}$ might then be achievable.

For a rocket to overcome Earth's gravity, it must have a thrust larger than its weight. Nuclear rockets tend to be heavy, with a lower thrust-to-weight ratio than chemical rockets. For this reason, plus the fact that their exhaust can contain radioactive isotopes, they are more suited for upper stages high in the Earth's atmosphere or for orbital transfers to other planets.

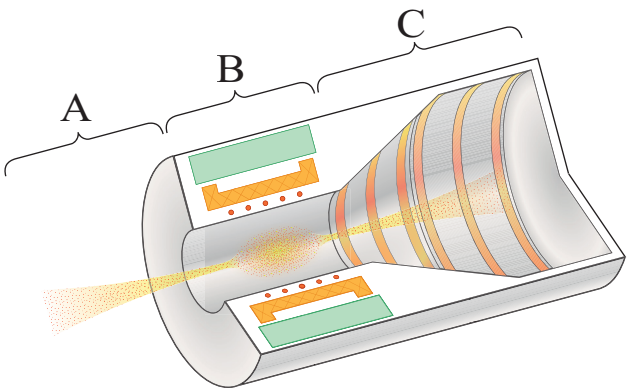
NASA's NERVA programme did a lot of development, tested several designs, and was not far from flight tests before the programme was cancelled for a combination of political and financial reasons, after having cost around 1.4 billion US dollars. There have since been further developments and plans to use such engines for missions to the Moon and Mars, but nothing concrete. As of mid 2025 no such thermal fission rocket has been publicly launched. However, it appears that the Russian military is developing a missile powered by such an engine, called the 9M730 Burevestnik (NATO name SSC-X-9 Skyfall). It was apparently tested in 2025, but details of its design and status remain unclear.

8.3.2 Fission fragment rockets

Fission reactions in the solid thermal nuclear engine only achieve propellant velocities of the order of $10\,\mathrm{km\,s^{-1}}$, just twice that of a chemical thermal rocket. The large energy density does allow a large amount of propellant to be heated, however, thus generating a large thrust. Liquid or gas cores could increase v_e, but it would still be at least two orders of magnitude below the theoretical limit of $11\,700\,\mathrm{km\,s^{-1}}$ achievable from uranium-235 fission.

The obvious way to attain this limit is to use the fission product nuclei directly as propellant. This is challenging because in the rest frame of the reaction, the product nuclei are moving in random directions. As the nuclei are charged, a magnetic field can

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009



A fission fragment rocket. The fission reactions take place in region B. The high-energy fission products are collimated by magnetic fields and propelled out of the back of the spacecraft in region A. Some of the product nuclei could be used within the spaceship to provide energy for onboard power (region C). Credit: Duckysmokton, licensed under CC-BY-2.5, from Wikimedia Commons, modified by the author; following the design of Clark and Sheldon (2005).

Fig. 8.5

be used to direct them into a collimated propellant beam, something we will examine in section 8.4.

Unlike in a nuclear reactor, where the fissile material is usually arranged into multiple thin rods, the fuel in a fission fragment rocket is in the form of small dust-sized grains (10 μm or less) to allow the fission fragments to escape from the fuel (figure 8.5). By ionizing the grains, electric fields can be used to contain them within the combustion chamber. As the fuel grains move slowly, they have a much lower kinetic energy to charge ratio than the fission products. The electric field can be tuned to contain the fuel without having much effect on the fission products, which then move off to be collimated by the magnetic nozzle and exhausted as high velocity propellant. At any one time there needs to be a critical mass of fuel dust in the reactor to sustain the reaction, so fuel must be continuously replenished to balance what is ejected as propellant. Control rods moderate the reaction rate as in a normal reactor. The positively charged fission fragments must be neutralized as they leave the spacecraft to avoid the spacecraft becoming negatively charged with time, as also done with the ion engine of chapter 6.

A specific design by Werka et al. (2012) uses a 1 GW reactor several metres in size with an engine mass of 110 t to produce an effective exhaust velocity of $v_e = 5170\,\mathrm{km\,s^{-1}}$, or $0.017\,c$. About half of the engine mass is the moderator, and much of the rest is the magnetic/electric field generators. The engine thrust is low, just 43 N, on account of the low mass flow rate of $\dot{m} = 8.3 \times 10^{-6}\,\mathrm{kg\,s^{-1}}$, or 0.72 kg per day. This corresponds to a power of $(1/2)\dot{m}v_e^2 = 110\,\mathrm{MW}$, so about 11% of the energy of the reactor is going into the kinetic energy of the propellant. Even though the fuel is presented in dust-sized grains to maximize the escape probability of the fission fragments, some fragments will inevitably collide with unburned fuel, thereby dissipating some of their kinetic energy as heat. A cooling system is therefore essential. Some of this heat could be captured and used to drive a generator in a closed cycle heat engine, thus generating electrical energy for the spacecraft. But the heat flux is too large for all of it to be captured in this way, so most has to be radiated into space as infrared radiation. The study authors calculate that 700 MW has to be radiated in this way, 70% of the reactor power.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

As the thrust of this fission fragment engine is very low and its mass large, the engine has to be operated for an extended period in order to achieve a large Δv. If we optimistically assume an initial spacecraft mass of 110 t (i.e. dominated by the engine; in practice the cooling system is more massive) and operate it for n days, then from the rocket equation (equation 3.14) the velocity achieved is

$$\Delta v = 5170 \ln(1.1 \times 10^5/[1.1 \times 10^5 - 0.72n])\,\mathrm{km\,s^{-1}}\,. \tag{8.1}$$

For $n = 1$ this is $0.034\,\mathrm{km\,s^{-1}}$; for $n = 365$ it is $12.3\,\mathrm{km\,s^{-1}}$. Even after 100 years the rocket is only moving at $1410\,\mathrm{km\,s^{-1}}$ ($0.0047\,c$). After this time it has consumed 26 t of fissile material (more or less equal to the ejected mass) and travelled 0.22 ly (equation 3.15b). If it cruised at constant velocity from this point on, the spacecraft would need another 860 years to travel the remaining distance to Alpha Cen. Even though v_e is high, because the mass ratio is low, it takes a long time to achieve the final Δv, which is also not very high. In reality the initial mass would be much larger to accommodate the propellant and spacecraft structure, so the mass ratio would be even lower and the travel time longer. This engine was not designed for an interstellar mission, but this calculation nonetheless shows that even such a large powerful fission fragment engine will not get us to the nearest stars very quickly.

The engine could be scaled up in power to try to shorten the mission length. Suppose we increase the power of the engine by a factor of 10 at constant v_e using the same fission reactions and efficiency (so the propellant power also increases by a factor of 10, to 1100 MW). This increases the mass flow rate and hence the thrust by a factor of 10. The 26 t of fuel will now be used up 10 times faster, within 10 years. After this time the spacecraft is moving at the same velocity as before (v_e and the mass ratio are unchanged) and it has moved a tenth as far (0.022 ly). Although it has achieved its cruising velocity earlier, the spacecraft now has a larger distance left to Alpha Cen, and in fact still needs 899 years to get there. The total travel time is 909 yr, not much less than the 960 years in the previous case, despite a 10-fold increase in power. The limiting factor is the amount of propellant, not how quickly we can expel it.

Suppose we instead double v_e to $10\,340\,\mathrm{km\,s^{-1}}$ ($0.035\,c$), which is only slightly less than the maximum possible with uranium-235. The spacecraft and propellant mass, as well as the mass flow rate, are unchanged, so the thrust doubles. It still takes 100 years to exhaust all the propellant, after which time the spacecraft is travelling at $2820\,\mathrm{km\,s^{-1}}$ ($0.0094\,c$) and has moved 0.45 ly. Cruising at this speed, the spacecraft needs another 404 years to get to Alpha Cen, for a total journey time of 504 years. As most of the journey is spent cruising at twice the speed as the nominal case, the journey time is nearly halved.

If we varied the propellant mass or accelerated for a significant part of the journey, then the scalings are not as straightforward as above. This is explored in an exercise at the end of the chapter.

Returning now to original engine design above, Werka et al. (2012) used it to outline a crewed return mission to Jupiter. In this, the total mass is around 300 t, a large part of which is $6000\,\mathrm{m^2}$ of radiators to remove the excess heat from the engine. This mass is about a third that of a comparable mission they outline using a thermal fission rocket (section 8.3.1), but the round trip time is 16 years as opposed to 4.5 years, on account of the much lower thrust.

The exhaust of a fission fragment engine is highly radioactive, so we would not want to operate such a rocket in the Earth's atmosphere. It would anyway be useless for getting into space, not only because of the low thrust, but also because the fission fragments would be

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

absorbed by the air before they could be magnetically coupled to the rocket. In practice such a rocket would be launched into space by conventional chemical rockets (perhaps in parts and then assembled there) and the fission engine only started once well away from the Earth.

Fissile materials other than uranium-235 could be used in rockets and may be preferred. Americium-242, for example, has a higher fission cross-section and produces more neutrons per fission, which lowers the critical mass. It does not occur naturally, however, so would have to be created by neutron capture in americium-241, which in turn is produced by beta decay of plutonium-239. Americium-242 has a half-life of 141 years, so for many missions it could be synthesized on the ground before launch.

8.3.3 Fission pulse rockets

A nuclear fission explosion occurs when a critical mass of fissionable material is brought together rapidly, such that the neutron capture rate becomes supercritical. If supercriticality can be retained for long enough, there will be a rapid release of large amounts of energy. The fission bomb dropped on Hiroshima by the USA in 1945 consisted of two pieces of uranium with a combined mass of 64 kg and a uranium-235 enrichment of 80%. When brought together they became supercritical, and this was retained long enough to release 15 kt TNT-equivalent of energy. A kiloton (kt) TNT-equivalent is a unit of energy equal to 4.184×10^{12} J often used to measure explosive strength. The bomb was very inefficient, as only about 1% of the uranium-235 underwent fission. The second bomb, dropped on Nagasaki, comprised two pieces of plutonium-239 with a combined mass of 6.2 kg. This is actually subcritical; the weapon used neutron reflectors to achieved supercriticality. This bomb used a shell arrangement (as opposed to a gun) leading to 20% of the plutonium undergoing fission during the explosion, releasing 21 kt TNT-equivalent of energy.

The concept of nuclear pulse propulsion is straightforward: explode a small nuclear bomb near a spacecraft and let the momentum of the fragments propel it. In its simplest incarnation the spacecraft is equipped with a 'pusher plate', which absorbs the momentum of an explosion that takes place outside the spacecraft. This plate is attached to the rest of the spacecraft by what is essentially a large spring, so that the impulses provided by a series of explosions can be smoothed out into a more continuous acceleration of the vehicle as a whole (figure 8.6). Provided the plate is massive enough and of the right materials, then a series of small explosions should individually not heat the plate enough to destroy it, and also allow enough time between explosions for it to cool. A thin layer at its surface will be vaporized by each explosion, but this could be built into the design. Effective exhaust velocities of 30–100 km s^{-1} could be achievable in this way. By using a magnetic field to couple the plate to the fission products of the explosion, the ablation at the surface could be reduced and more powerful explosions used, increasing v_e to perhaps 1000 km s^{-1}.

A variation on this idea is for the explosions to occur within a combustion chamber inside the rocket. The motivation is to transfer more of the momentum from the explosion by capturing also the fragments that were moving away from the spacecraft and not intercepting the pusher plate. This could be done by using the explosion to heat up a propellant, such as hydrogen, which is then ejected through a nozzle just as with any other thermal rocket. This would again be pulsed, with propellant injected into the combustion chamber for each explosion. This approach was deemed infeasible, however. Part of the problem is that some of the energy from the nuclear explosions is carried by neutrons. These mostly pass through

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

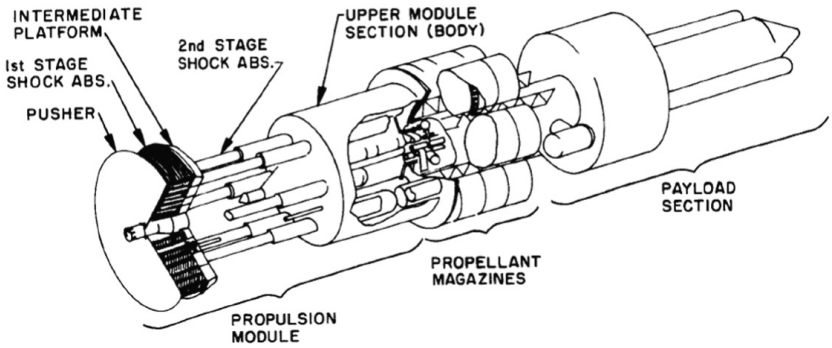


Fig. 8.6 The design of a pulse fission rocket produced as part of the NASA project Orion in the 1960s. The pusher plate at the left end has a diameter of about 10 m. The propulsion module has a mass of 100 t and was expected to deliver an effective exhaust velocity of about 20 km s^{-1}. Credit: NASA, public domain, from Wikimedia Commons.

the propellant but are then absorbed by the walls of the combustion chamber, heating that instead. This would require significant cooling. The cooling apparatus as well as the combustion chamber itself add significant mass. Theoretical studies indicated that the effective exhaust velocities achievable by such internal engines would be an order of magnitude lower than the external, pusher-plate arrangement.

The idea of nuclear pulse propulsion was explored experimentally in the USA by the Department of Defence and NASA between 1958 and 1965 under the name Project Orion (figure 8.6). It went through a series of designs, including ones launched from the Earth's surface, which would of course have produced significant amounts of nuclear fallout. Each bomb would release 0.01–0.1 kt TNT-equivalent of energy. One design for a crewed round-trip mission to Mars was to be started from Earth's orbit. It used a 100 t engine that provided an effective exhaust velocity of 20 km s^{-1} and a thrust of 3.5 MN. The total spacecraft mass was 1100 t, of which 150 t was payload. Using one to two thousand fission explosions, the travel time to Mars was to be 125 days. The spacecraft included shielding to reduce the radiation dose for the crew from the engines down to 500 mSv for the round-trip mission, the same as the additional dose expected from cosmic rays.[3]

One of the experiments undertaken for Project Orion involved exploding a 20 kt TNT-equivalent fission bomb 10 m away from two 1-m-diameter graphite spheres. Surprisingly they survived, suffering just the ablation of a fraction of a millimetre of their surfaces. Before testing nuclear bombs for propulsion itself, tests were done using chemical explosions. The most successful lifted an object 56 m into the air in 23 s using a series of seven explosions. Towards the end of the project, preliminary designs were being made for a spacecraft that would be assembled in Earth's orbit and launched from there. Ultimately, however, the project was cancelled in the wake of the 1963 nuclear test ban treaty, which banned the testing of nuclear weapons in the atmosphere, under water, or in space (underground tests were still permitted).

[3] The Sievert measures the biological effect of ionizing radiation, whereby 1 Sv = 1 J kg^{-1}. The average annual dose that a person receives on the Earth is of order 2.5 mSv, although it varies with location and activity. A chest X-ray delivers around 0.1 mSv and a whole body CT scan around 10 mSv, for example, while a 12-hour commercial flight delivers a dose of around 0.03 mSv. NASA and ESA limit an astronaut's lifetime exposure to 1000 mSv. The health effects of radiation also depend on how quickly the dose is received.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

The idea of using small bombs to propel an interstellar spacecraft was not dead, however, as we shall see in section 8.6.

8.4 Magnetic nozzles

A key idea underlying several types of nuclear propulsion is to couple the charged products to the rocket using a magnetic field. How does this work?

Consider a particle moving at velocity v relative to a rocket. If the particle is not coupled to the rocket, it moves off at this velocity and the rocket remains at rest. If the particle is instead coupled to the rocket, it still moves at velocity v relative to the rocket, so there has been no momentum or energy transfer in the rocket's instantaneous rest frame. But due to conservation of momentum, the rocket is now moving in the initial rest frame, so in this frame the particle must be moving in the opposite direction at a velocity less than v. In this frame the particle has transferred momentum and kinetic energy to the rocket. This is of course the basis of the rocket equation (section 3.1). Conceptually it does not matter how this coupling takes place. It could be via the elastic reflection of a gas molecule, or via a magnetic field for a charged particle. Just as a magnetic field accelerates a charged particle, so a moving particle will accelerate whatever the magnetic field is attached to, in this case the rocket.

Immediately after a set of nuclear reactions, there are numerous charged particles moving in random directions. These have zero net momentum in their centre-of-mass frame, which is also the rocket's instantaneous rest frame. We want to convert these random motions – thermal energy – into directed kinetic energy. Recall that this is the purpose of the convergent–divergent nozzle in thermal rockets (chapter 5). By collimating the particle motions, more momentum is transferred along the axis of the rocket rather than perpendicular to it.

In chemical rockets, the molecules of the hot gas transfer momentum by bouncing off the inside walls of the nozzle. With nuclear reactions, the temperature of the products is so high it would melt the rocket. We can instead use a magnetic field to couple the charged particles to the spacecraft without them making physical contact. The magnetic field is designed to convert the radial motions (perpendicular to the axis) into axial ones, and just like the convergent–divergent nozzle this converts thermal energy into directed kinetic energy. This is the *magnetic nozzle*.

How strong does the magnetic field need to be? We can get an order-of-magnitude estimate by computing the magnetic field strength required to keep a particle moving in a circle of fixed radius. To do this we equate the Lorentz force, qvB, with the (relativistic) centripetal force, $\gamma mv^2/r$, to give

$$B = \frac{\gamma mv}{qr} . \tag{8.2}$$

The radius r is sometimes called the Larmor radius. In a uranium-235 fission fragment rocket (section 8.3.2), an example product nucleus is $^{139}_{56}$Ba with a kinetic energy of 112 MeV. This corresponds to $v = 0.042\,c$, $\gamma \simeq 1$ (not very relativistic), and so a field strength of $B = 0.33$ T is required to turn it in a circle of radius 1 m. This is a strong field that would need large massive magnets to collimate the propellant.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

The actual field configuration required for the magnetic nozzle is much more complicated than this simple calculation implies. In particular, the field must diverge in the right way at the nozzle exit to ensure the plasma detaches from the rocket.

8.5 Nuclear fusion

When two nuclei with low atomic mass numbers combine, their product often has a higher binding energy per nucleon than the parent nuclei (figure 8.1). This combination, or fusion, releases the binding energy as the kinetic energy of the fusion product and other released particles.

The difficulty with achieving fusion is that nuclei are positively charged and so repel each other. For fusion to occur, the nuclei need to have sufficient kinetic energy to overcome this Coulomb barrier (or actually somewhat less, due to quantum tunnelling). This energy can be supplied by heating the fuel to high enough temperatures, of the order of tens or hundreds of millions of kelvins, which also turns it into a fully ionized plasma.

Fusion takes place in the centres of stars in the middle part of their lives and is the source of their luminosity. In lower-mass stars like the Sun (core temperature 1.6×10^7 K), the dominant reaction is the fusion of hydrogen (H, $Z = 1$) into helium (He, $Z = 2$). This takes place through a series of steps starting with[4]

$$\mathrm{p} + \mathrm{p} \;\rightarrow\; \mathrm{D} + \mathrm{e}^+ + \nu_\mathrm{e} + 0.42\,\mathrm{MeV}. \tag{8.3}$$

The deuteron and another proton then undergo another fusion to produce a helium-3 nucleus plus a gamma photon. Two such helium-3 nuclei then combine in yet another fusion reaction to produce a helium-4 nucleus (an alpha particle) plus two protons. The overall reaction is

$$6\mathrm{p} \;\rightarrow\; {}^4\mathrm{He} + 2\mathrm{p} + 2\mathrm{e}^+ + \nu_\mathrm{e} + 2\gamma \;+ 24.68\,\mathrm{MeV}. \tag{8.4}$$

Each positron quickly annihilates with an electron in the plasma, in each case producing two gamma photons (with a combined energy of 1.02 MeV), so the total energy released by the reaction is 26.72 MeV. This appears as the kinetic energy of the protons and alpha particles as well as the gamma photons and neutrinos. The fraction of the initial rest energy released[5] is $26.72\,\mathrm{MeV}/(4m_pc^2) = 0.0071$, where m_p is the mass of the proton. Note that the initial rest energy is considered to be four protons, not six, because the two protons produced can be used again in this reaction.

The above process is one of three branches of a more complicated reaction known as the proton–proton (p–p) chain, other branches involving the intermediate production of heavier elements. The details are unimportant from the point of view of propulsion systems, however, for one simple reason: the reaction rate of two protons is extremely small, a consequence of it being mediated by the weak interaction (something we can conclude from the production of a neutrino). In the Sun, for example, the average time before a proton fuses with another

[4] I use the symbols $\mathrm{D}={}^2_1\mathrm{H}$ and $\mathrm{T}={}^3_1\mathrm{H}$ for deuterium and tritium, the two heavy forms of hydrogen, respectively. p and n indicate the proton (${}^1_1\mathrm{H}$) and neutron (${}^1_0\mathrm{n}$) respectively. e^+ is the positron and ν_e is the electron neutrino.

[5] Conventionally this fraction is denoted ϵ. In the context of rockets, where some of the particles produced cannot be used for propulsion, the energy useable for propulsion is smaller, and was denoted by the factor $\eta\epsilon$ in figure 7.9.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

proton is nine billion years, which not coincidentally is close to the main sequence lifetime of the Sun.

We must turn to other fusion processes for propulsion. One of the most important from the point of view of future energy production on the Earth is the fusion of deuterium and tritium

$$\mathrm{D} + \mathrm{T} \rightarrow {}^4\mathrm{He}\ (3.52\,\mathrm{MeV}) + \mathrm{n}\ (14.07\,\mathrm{MeV}). \tag{8.5}$$

The energy released equals the difference in the masses, 17.59 MeV. The fraction of the initial mass converted into energy is 0.0038. This goes entirely into the kinetic energies of the product nuclei, in a proportion that can be determined here by the conservation of momentum. The momenta of the two products in the initial centre-of-mass frame must be equal in magnitude, which in Newtonian mechanics is $p = \sqrt{2Km}$. Hence the kinetic energy K of each product is inversely proportional to its mass m. The helium-4 nucleus therefore receives 3.52 MeV kinetic energy and the neutron four times as much, 14.07 MeV. This is a Newtonian calculation, but as the velocity of the neutron in the centre-of-mass frame is only about $0.17\,c$, it is a reasonable approximation. In reactions where there are three or more products, conservation of energy and mass does not yield a unique solution to the energy partition, and a range is possible (as in equations 8.3 and 8.4).

The above D+T fusion reaction produces 3.5 MeV per nucleon, compared to 0.8 MeV per nucleon for the fission of uranium-235. This is because the D+T reaction produces a nucleus – the alpha particle – that is unusually tightly bound for its atomic number: Note the peak of the binding energy per nucleon at helium-4 in figure 8.1.

In order to overcome the Coulomb barrier and get a significant probability of fusion occurring, the nuclei require a combined energy of order 10 keV. This corresponds to a temperature of $T = 8 \times 10^7$ K, from $(3/2)kT = E$, where k is the Boltzmann constant. (Recall that the temperature only represents the average kinetic energy of the the plasma. Fusion can occur in a cooler plasma in the high-energy tail of the energy distribution.) This kinetic energy is nonetheless much smaller than the energy released, so can be neglected in the energy partition just mentioned.

The D+T reaction is being intensely pursued for large-scale energy production on the Earth. Deuterium is a stable, naturally occurring isotope of hydrogen. Although it only makes up a small proportion of hydrogen on the Earth, just 0.016% by number, hydrogen is widely abundant in water, and its extraction is routine. Deuterium has been used for many years as a moderator in fission reactors that use natural (non-enriched) uranium (section 8.2). Tritium, on the other hand, is a radioactive isotope with a half-life of just 12.3 years. It barely occurs naturally on the Earth (mostly via the interaction of cosmic rays with the atmosphere), and unnaturally as a result of nuclear weapons testing. Tritium must therefore be manufactured, usually via the reaction

$${}^6\mathrm{Li} + \mathrm{n} \rightarrow {}^4\mathrm{He}\ (2.06\,\mathrm{MeV}) + \mathrm{T}\ (2.75\,\mathrm{MeV}) \tag{8.6}$$

by lining a nuclear reactor with lithium (Li, $Z = 3$) to take advantage of the energetic neutrons produced therein. That is, the tritium fuel required for the fusion (equation 8.5) can be bred within the reactor itself. Lithium-6 is a naturally occurring stable isotope of lithium (1.9–7.8% by number).

The probability of a fusion reaction occurring is determined by the cross-section (its unit is area). A significant probability is only achieved once the temperature reaches tens of millions of kelvins, as can be seen in figure 8.7 for several different reactions. The plasma must also

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

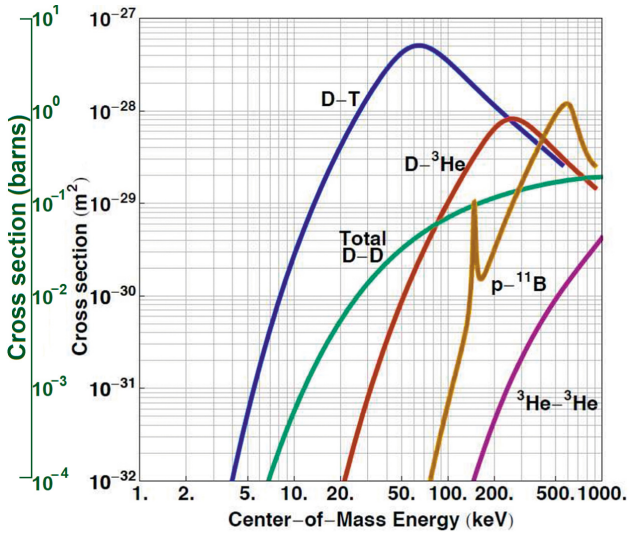


Fig. 8.7 The cross-section for different fusion reactions as a function of the kinetic energy (temperature). An energy of 10 keV corresponds to a temperature of 8×10^7 K. Credit: WikiHelper2221, licensed under CC-BY-SA-4.0, obtained from Wikipedia Commons.

be dense enough to get a significant amount of fusion. Such a hot plasma will lose energy, the rate of which is measured by the inverse of the confinement time. In order for a fusion reaction to be self-sustaining, the amount of energy produced needs to exceed the energy losses so that the temperature can be maintained. These three aspects indicate that for significant fusion to occur, there is a minimum of the product of the ion number density n, the temperature T, and the confinement time τ. This triple product, sometimes also referred to as the *Lawson criterion*, is often used as a figure-of-merit of fusion systems. For D+T fusion to produce power in a reactor, a triple product $nT\tau$ in excess of a few $10^{21}\,\mathrm{m}^{-3}\,\mathrm{s\,keV}$ must be achieved.

No solid material can tolerate the high temperatures required for fusion, so much research has been done on two different approaches to contain the plasma within a small volume. The first approach is to heat the reactants – which ionize to become a plasma – with an electric current via Joule heating. Strong magnetic fields are then used to keep the nuclei at high enough densities for reactions to occur, giving this approach its name, *magnetic confinement fusion*. One type of reactor, known as a tokamak, is key to the International Thermonuclear Experimental Reactor (ITER). The plasma will be heated to $1\text{–}1.5\times10^8$ K and retained at this temperature for extended periods while the plasma is confined to a small volume using powerful superconducting magnets that produce a field strength of the order of 10 T. With a plasma density of around $10^{20}\,\mathrm{m}^{-3}$ and confinement times of a few seconds, this will enable a self-sustained net-power-producing reaction. ITER is a large reactor, several metres in radius with a plasma volume of nearly $1000\,\mathrm{m}^3$, intended to provide around 500 MW of thermal power. Most of the energy from the D+T reaction in equation 8.5 is in the neutron, so to generate electricity in a reactor (not planned for ITER), the neutrons as well as the alpha particles need to be absorbed to heat a working fluid. This can then be used to drive a high-pressure steam turbine to generate electricity, just like in other power stations.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

The second approach to achieving fusion goes by the name of *inertial confinement fusion*. The approach here is to manufacture the fuel into mm-sized pellets and then to use high-power lasers to compress and heat these pellets to achieve fusion, releasing the energy in a single explosive reaction. One such experiment, the National Ignition Facility (NIF) in the USA, uses 192 lasers with a combined power of 500 TW, although these only operate in bursts lasting a few nanoseconds. The process is not very efficient, with only about 0.5% of the initial 2 MJ of energy from the laser light ending up in the pellets. However, this is enough to compress them to very high densities (at pressures of 10^{16} Pa) and heat them to about 10^8 K, causing the deuterium and tritium in the pellets to undergo fusion. Theoretically this could yield in the order of 100 MJ per pellet. By the end of 2022 only about 3 MJ could be produced, but this was above the 2 MJ of laser energy put in, making it the first fusion experiment to achieve an energy gain Q above 1.0. Although a significant milestone, this overlooks the 400 MJ of energy that was required to charge the laser. NIF is a science experiment. Clearly, to become a useful source of domestic power, or to form the basis of a propellant system, pellets must be fused in rapid succession, but currently the turnover time is many hours.

Magnetic confinement fusion needs large superconducting magnets, which are both massive and consume a lot of energy, making them rather unsuitable for spacecraft. ITER, for example, will have an external diameter of 20 m and a mass of around 5000 t. The lasers of inertial confinement fusion are likewise massive and energy-hungry (just the NIF target chamber is 10 m in diameter and weighs 130 t). Thus the facilities currently being developed for experiments and power plants cannot be used in practical propulsions systems. However, some of the ideas and results coming from this work have inspired the development of hybrid systems that may ultimately lead to practical propulsion methods, as we shall see in section 8.7.

8.5.1 Reactions for rockets

While the D+T reaction is suitable for energy generation on the Earth, it is not suitable as a basis for propulsion for one main reason: 80% of the kinetic energy is carried by the neutron. Being neutral, its momentum cannot be coupled to the spacecraft using a magnetic field (section 8.4), so this energy is essentially lost. With the D+T reaction we could only really make use of one-fifth of the energy released, namely the 3.52 MeV in the alpha particle.

In practice, some of the neutrons from the D+T reaction would collide with the spacecraft and be stopped by it, thereby transferring some of their kinetic energy to the spacecraft. But this is actually a problem for a number of reasons. First, the collisions are inelastic, so release a significant amount of energy that heats the spacecraft. The spacecraft then needs large radiators to remove this heat to avoid melting. Second, these high-energy neutrons can create collision cascades, thereby dislocating atoms from lattices to create point defects that degrade the spacecraft structure. Third, some neutrons would be captured by nuclei in the material of the spacecraft, creating excited, unstable nuclei. These then decay radioactively, releasing alpha particles, electrons, and/or gamma photons, or even undergoing fission, a process known as neutron activation. This is unlikely to be good for the structural integrity of the spacecraft and would demand the use of shielding to protect its instruments. We therefore look to other reactions.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

A more promising reaction is the fusion of deuterium with helium-3:

$$\mathrm{D} + {}^3\mathrm{He} \rightarrow {}^4\mathrm{He}\ (3.67\,\mathrm{MeV}) + \mathrm{p}\ (14.68\,\mathrm{MeV})\ . \tag{8.7}$$

This too produces a helium-4 nucleus and releases a similar amount of energy as the D+T reaction. But instead of a neutron, it produces a proton. Being positively charged, its energy can be coupled to the spacecraft electromagnetically, potentially allowing us to harness all of the energy released (but see below).

This reaction presents two difficulties compared to D+T, however. The first is that because there are two protons in helium-3 as opposed to just one in tritium, the Coulomb barrier is higher. Energies of order 60 keV must be attained in order to get significant fusion, several times higher than that required for D+T (see figure 8.7). The second difficulty is that helium-3 hardly exists on the Earth, even though it is a stable isotope. It is a primordial isotope, making up 1–10 ppm of the helium on the Earth (which itself is very rare; only 5 ppm of the atmosphere), but most is in the Earth's mantle and so inaccessible. Helium-3 does occur in natural gas sources, but at current rates of gas extraction only a few kg per year could be collected in this way. It is also produced in small quantities by some reactions, including the radioactive decay of tritium, itself produced by reaction 8.6, for example, meaning helium-3 is a trace product from nuclear power stations and weapons testing. Some is also produced by cosmic ray spallation. But these sources are far too small to be a source of fuel for future rockets.

There may be more helium-3 on the surface of the Moon, as the result of interactions of the solar wind with lunar regolith. Estimates are of a few parts per billion in lunar surface rock, a thousand times more than in the Earth's atmosphere. Helium-3 appears to be more abundant in the outer atmospheres of the gas giant planets. The Galileo probe, for example, measured the ${}^3\mathrm{He}/{}^4\mathrm{He}$ ratio in the atmosphere of Jupiter to be 1.7×10^{-4}. Taking the abundance of helium to be 0.14 by number, this corresponds to a helium-3 number density of around 20 ppm. Although mining helium-3 from there may seem like a stretch of the imagination, not least because of Jupiter's deep gravitational potential well, that may not be the hardest aspect of an interstellar mission.

Another reaction that could be suitable for rockets is the fusion of two deuterons. This can be achieved with a slightly lower temperature than D+^{3}He (figure 8.7), of the order of 30 keV, and has two branches that occur with equal probabilities:

$$\mathrm{D} + \mathrm{D} \rightarrow \mathrm{T}\ (1.01\,\mathrm{MeV}) + \mathrm{p}\ (3.02\,\mathrm{MeV}) \quad 50\% \tag{8.8a}$$

$$\mathrm{D} + \mathrm{D} \rightarrow {}^3\mathrm{He}\ (0.82\,\mathrm{MeV}) + \mathrm{n}\ (2.45\,\mathrm{MeV}) \quad 50\%. \tag{8.8b}$$

The first reaction produces two charged products. The second, in contrast, gives three-quarters of its energy to the neutron. These reactions produce less useable energy than the D+^{3}He reaction, with only 66% of the energy going into the charged particles. However, the products in both branches can undergo further fusion reactions. If the conditions are right, the tritium in the first branch can fuse with another deuteron via the D+T reaction (equation 8.5), and the helium-3 in the second branch can fuse with another deuteron via the D+^{3}He reaction (equation 8.7). The overall reaction is

$$3\mathrm{D} \rightarrow {}^4\mathrm{He} + \mathrm{p} + \mathrm{n} + 21.62\,\mathrm{MeV}. \tag{8.9}$$

This is sometimes referred to as the catalyzed deuterium reaction. The energy partition among the products is different for the two branches, but overall 62% of the energy goes into the charged particles, and the rest into the neutrons.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

The catalyzed deuterium reaction puts more energy in the charged particles per input nucleon than the D+^{3}He reaction (4.76 MeV vs 3.67 MeV), implying a better energy-to-fuel mass ratio, and it uses a readily available fuel, just deuterium. The D+^{3}He reaction appears to have the advantage that it does not produce any neutrons, which can damage the spacecraft through the process of neutron activation mentioned earlier. However, while this is true for the principal reaction, there are inevitably side reactions that do create neutrons, in particular equation 8.8b. Equation 8.8a creates tritium, which can either react with deuterium via equation 8.5 or react with itself, both reactions creating neutrons. Thus the production of neutrons, and the problem this creates, is more or less unavoidable. In practice, though, only about 5% of the kinetic energy would be lost to neutrons resulting from D+D. The lower neutron flux from D+^{3}He gives it a considerable advantage over the catalyzed deuterium reaction.

Several more side reactions occur when trying to fuse D+^{3}He (or indeed other nuclei), not only between D and D, D and ^{3}He, and ^{3}He and ^{3}He, but also between these and the products of those reactions, which include hydrogen, tritium, and small quantities of lithium-5 and beryllium-6. Some reactions give off photons too. Which reactions take place with what relative frequencies depends on the exact conditions of the fusion, in particular the temperature. Thus in practice the products of fusion, and the energy that can be extracted from them, is more complex than has been summarized here.

The fuel for any practical fusion reaction would be stored in atomic or molecular form, and this fully ionized into a plasma before the above fusion reactions occur. But the ionized electrons have not disappeared. Through collisions with the nuclei, the electrons become heated, gaining kinetic energy at the expense of the nuclei. These moving electrons are then decelerated by the electric field in the plasma (generated by the positive nuclei), and so lose some of their kinetic energy again in the form of X-ray photons, a process known as *bremsstrahlung*. If the plasma is not optically thick, these photons will remove energy from the plasma. Bremsstrahlung is an important source of energy loss from fusion reactions, especially for the D+^{3}He and D+D reactions. As this reduces the kinetic energy of the charged fusion products, it also reduces the effective exhaust velocity of a fusion-powered rocket.

Bremsstrahlung X-rays are furthermore problematic as a source of ionizing radiation that can damage the spacecraft structure and as a source of heat that can melt the spacecraft. It may be possible to absorb some of the bremsstrahlung in a fluid that is used in a closed-cycle heat engine to generate electricity. Better still is a direct conversion of the heat to electric power that avoids moving parts, for example via the thermoelectric effect. This could be used either to power systems onboard the spacecraft, or it could be used to drive an electric field that accelerates the charged fusion products further (as in an ion engine; chapter 6). In practice the spacecraft would still need large radiators to remove the excess heat, and these add significant additional mass.

There exist many other fusion reactions, but most are less practical because they have lower fusion cross-sections, require higher temperatures to initiate fusion, or involve nuclei with higher atomic numbers. The latter is problematic because the intensity of the bremsstrahlung is proportional to the square of the atomic charge.

One final reaction worth mentioning, however, is the interaction[6] of boron (B, $Z = 5$) with a proton to produce three alpha particles,

$$^{11}\mathrm{B} + \mathrm{p} \rightarrow 3\,^4\mathrm{He} + 8.7\,\mathrm{MeV}, \tag{8.10}$$

[6] This reaction is actually fission, not fusion, but it is usually considered among fusion reactions.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

because this uses a readily available fuel. The energy released per nucleon is relatively low compared to the earlier reactions examined because boron-11 has a large binding energy per nucleon. It also needs a very high temperature to achieve a reasonable cross-section (several hundred keV; see figure 8.7). Although nominally aneutronic, side reactions such as boron reacting with an alpha particle to produce nitrogen and a neutron can remove some of the kinetic energy. However, only about 0.2% of the total energy yield is taken away by neutrons. It does, however, produce significant amounts of bremsstrahlung, because the fully ionized boron fuel contributes a lot of electrons.

8.5.2 Effective exhaust velocities

When the combined kinetic energy of the fusion products is K, and the sum of their masses is m, then classically we can write an equivalent velocity as $\sqrt{2K/m}$. Assuming no energy losses, this is the effective exhaust velocity of the propellant. The quantities K and m only include the particles that can be coupled to the rocket, so this usually excludes the neutrons. For the D+^{3}He reaction in equation 8.7, $K = 18.35\,\text{MeV}$ and $m = 5.009\,m_\text{u}$, so $v_e = 0.0887\,c$. This is well below the speed of light, so this classical calculation is a reasonable approximation. The temperature corresponding to the kinetic energy is $2E/3k = 1.4\times10^{11}$ K, making it obvious that a physical nozzle could not be used to collimate this propellant.

An alternative (and relativistic) way to compute v_e is the approach in section 7.8. For this we need to compute the factors ϵ, η, and ζ. The energy released is the difference between the initial and final masses of the nuclei. For the D+^{3}He reaction, the sum of the initial masses is $2.013553 + 3.014932\,m_\text{u}$ or 4684.005 MeV, and the sum of the final masses is $4.001506 + 1.007276\,m_\text{u}$ or 4665.652 MeV, giving a difference of $0.0197\,m_\text{u}$ or 18.35 MeV. As the kinetic energy required by the parent nuclei to overcome the Coulomb barrier is only a few tens of kiloelectronvolts, it can be neglected. Assuming there are no losses ($\eta = 0$), then the factor ϵ is the fraction of the initial mass converted into energy, which here is $\epsilon = 18.35/4684.005 = 0.003918$. No photons are produced (matter rocket) so $\zeta = 1$. Equation 7.40 gives $v_e = \sqrt{\epsilon(2-\epsilon)} = 0.0884\,c$. This is very close to what we just computed using classical mechanics in this case.

8.6 Fusion pulse rockets

Having established that fusion reactions produce charged particles with very large kinetic energies, the question is how to use these to propel a rocket. The direct approach is to use a magnetic nozzle to couple the charged fusion products to the rocket, as explained in section 8.4.

One way of achieving this is pulse fusion, in which fusion takes place in sequential pulses, each giving the rocket a small impulse. A well-developed concept for this was presented in the 1970s by a team associated with the British Interplanetary Society. The mission, based on the D+^{3}He reaction, was called Daedalus after the mythological Greek character who made artificial wings to enable him to fly. The goal of the modern Daedalus was to carry an uncrewed payload with a mass of 450 t to Barnard's Star in 50 years. At 6.0 ly distance, this is the second closest stellar system to the Sun (section 2.5), and was chosen because at the

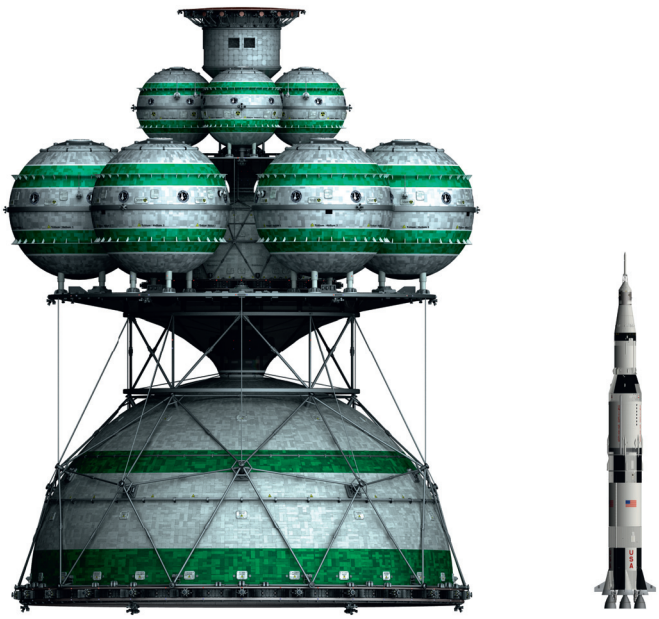
Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

time there were claims of the detection of a planet around the star. This journey requires the spacecraft to achieve an average velocity of $0.12\,c$.

According to one plan, Daedalus is fuelled by deuterium and helium-3 mixed into small pellets of mass 3 g (first stage) and 0.3 g (second stage). These are compressed by electron beams or lasers to initiate fusion, and assumed to provide an effective exhaust velocity of $10^4\,\mathrm{km\,s^{-1}}$, or $0.033\,c$. This is a lot lower than the ideal value of $0.089\,c$ derived in section 8.5.2 because of the inevitable inefficiencies in extracting kinetic energy from the fusion.

The total mass of the spacecraft is 52 670 t (figure 8.8) of which 50 000 t is fuel and the remainder is the structure (engines, tanks, payload, etc.) This is split across two stages. The first has a structural mass of 1690 t and 46 000 t of fuel. The second has a structural mass of 980 t, of which 450 t is the payload, plus 4000 t of fuel. The mass ratio (initial/final) of the first stage is $R = 52\,670/(52\,670 - 46\,000) = 7.90$ and thus delivers $\Delta v = 2.1 \times 10^4\,\mathrm{km\,s^{-1}} = 0.069\,c$ (assuming $v_e = 10^4\,\mathrm{km\,s^{-1}}$; the actual design study used slightly different values for the two stages). The structure of the first stage is then ejected. The mass ratio of the second stage is $R = (4000 + 980)/980 = 5.1$ and thus $\Delta v = 1.6 \times 10^4\,\mathrm{km\,s^{-1}} = 0.054\,c$. The total Δv is (classically) the sum of these, $3.7 \times 10^4\,\mathrm{km\,s^{-1}}$ or $0.123\,c$. The actual mission design was more optimized than this, with the fuel in each stage stored in multiple tanks, each of which is discarded as soon as it is empty, which slightly increases the Δv achievable with each stage.

The engines are designed to burn for an appreciable fraction of the mission time. In the first stage, pellets of mass 3 g are fused at a rate of 250 per second, each pellet releasing 170 GJ of energy (over four orders of magnitude more than is currently achieved per pellet



An artist's impression of Daedalus, 190 m tall, shown in comparison to a Saturn V rocket (111 m tall). Credit: Nick Stevens, http://nick-stevens.com/. Fig. 8.8

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

in the NIF), corresponding to a power of 44 TW. With 46 000 t of fuel this takes 2.05 years to exhaust. Lower-mass pellets are used in the second stage (each releasing 13 GJ), but fused at the same rate, taking 1.76 years to exhaust. A total of 30 billion pellets are consumed. From these mass expulsion rates we can also compute that the thrusts of the two stages are 7.1 MN and 0.72 MN, respectively.

The enormous propellant mass arises from the large payload mass of 450 t (via the rocket equation). This payload would permit a substantial amount of scientific equipment to be taken. However, this is a flyby mission with an encounter velocity of 0.9 au per hour, leaving little time to gather data near the target star. Daedalus addresses this by including 18 smaller spacecraft that are released well before the encounter. Although they would not decelerate in the nominal mission, they would each fly much closer to different targets of interest. They would send their data back to Earth via the main spacecraft, which repurposed its 40-m-diameter engine reaction chamber as a radio antenna.

It is historically interesting that the first planned scientific investigation to be done from Daedalus as it approached the target was to use onboard telescopes to look for planets. This was the late 1970s, and within 20 years the first exoplanets had been discovered around a nearby star. Given the continuing improvements in technology, and the strong focus on exoplanets by astronomers, we can reasonably expect all but the smallest planets around the target star to be discovered from the Earth (or nearby space telescopes) well before a mission would ever get to its target. But there would still be much to discover and explore from the spacecraft, as discussed in chapter 2.

Given the enormous mass of Daedalus, it was not intended for launch from the Earth's surface. Indeed, most of its mass is helium-3, which the study writers assumed would be mined from the Moon or from the atmospheres of the gas giant planets. The spacecraft would therefore be assembled in orbit about the source planet/Moon and launched from there. Exactly how helium-3 would be mined remained an open question.

We saw in section 8.3.3 that a lot of effort and money was investigated in the past into the use of nuclear fission bombs to power a rocket. Daedalus is similarly propelled by a series of explosions. One of the main differences, other than the magnetic coupling, is that with fusion there is no minimum mass for the reaction, so very small explosions can be used.

There are many other concept studies for interstellar spacecraft based on pulsed fusion engines, in particular using inertial confinement fusion. One set of studies followed up directly on Daedalus, going by the name of Icarus, who in Greek mythology is the son of Daedalus (although unlike his father, he did not survive their joint attempt to escape from Crete using artificial wings).[7] Long et al. (2011) optimized the Daedalus design for a fixed mass of propellant, investigating solutions with one to four stages, still as a flyby mission. They came to the preliminary conclusion that more stages were not necessarily better, because with a fixed propellant mass the extra stages add structural mass and thus decrease the payload mass for the given total mass.

Other Icarus designs focused on smaller payloads and so required much less propellant. One design has a total mass of 2700 t, of which 320 t is the engine, 150 t the payload, and 2000 t the D+^{3}He fuel. It is a single stage rocket (one engine with a jet power of 1.2 TW), but ejects fuel tanks once they are empty. After 2.2 years it would reach its cruise velocity of

[7] This should not be confused with a Soviet Daedalus-like project also called Icarus that appeared around the same time as Daedalus, but about which little has been published.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

0.045 c and it would take a total of 98 years to travel to Alpha Cen. This relaxing of the travel time requirement (from 50 years) permits smaller engine powers and smaller mass ratios, and so makes the mission far more doable.

Another variation of the Icarus design called Firefly uses a different plasma confinement approach called the z-pinch. When electrical currents are passed in the same direction along parallel wires, they generate magnetic fields that interact with the currents, creating Lorentz forces that attract the wires to one another. In the same way, a current run through a plasma will compress or 'pinch' the plasma, until the Lorentz force is balanced by the gas pressure. If the compression is large enough, it could initiate fusion. The Firefly design requires a large current of 5×10^6 A and a huge power to sustain this. It uses readily available deuterium in a D+D reaction, thus avoiding the need to find helium-3, but pays the price in having to cope with a large neutron flux that would overheat the spacecraft. Thus a significant fraction of the spacecraft's dry mass has to be dedicated to cooling, in terms of coolant fluid (hundreds of tonnes of beryllium), coolant pumps, and large radiators.

There are many variations of the Icarus design using different fusion reactions and approaches to fusion. They, like the original Daedalus study, demonstrate the complexities and trade-offs in designing a mission. What they have in common is very massive spacecraft of thousands of tonnes with payloads of tens of tonnes getting to one of the nearest stars within a hundred years.

8.7 Other fusion propulsion schemes

Several other approaches have been proposed that use fusion for rocket propulsion. Like the fusion engines of the previous sections, none of these systems have been developed, nor have they achieved net fusion power output. Although some informative experiments have been done with plasmas, many of the detailed results rely on computer simulations. Here we take a brief look at a few different ideas.

The first idea is to use a sustained fusion reaction. By siphoning off some of the high-energy fusion products from the plasma, and directing them with a magnetic nozzle, we obtain thrust. This is analogous to the fission fragment engine (section 8.3.2). One way to implement this is with mirror magnetic confinement, in which the plasma is trapped between two magnetic fields, one of which leaks a small amount of plasma into the nozzle. Fuel must be continuously supplied to the engine, and a sufficiently high temperature and density must be maintained such that the energy siphoned off for propulsion does not stop the reaction (recall the triple product). The challenge with this approach is that the mass flow rate of the products is low, and so despite the high velocity of the fusion products, the thrust remains low.

Another idea is represented by the *direct fusion drive* (DFD), in which energy from the fusion is used to heat up a separate propellant, in analogy to the thermal fission rocket (section 8.3.1). The DFD uses the D+^{3}He reaction of equation 8.7 and is based on the field-reversed configuration. This is a way of confining the plasma in a torus with a magnetic field (the tokamak and stellarator are other examples). A rotating magnetic field heats the plasma using radio-frequency waves, aiming to achieve a temperature of around 100 keV ($T \sim 10^9$ K) at a density of 5×10^{20} m^{-3}.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

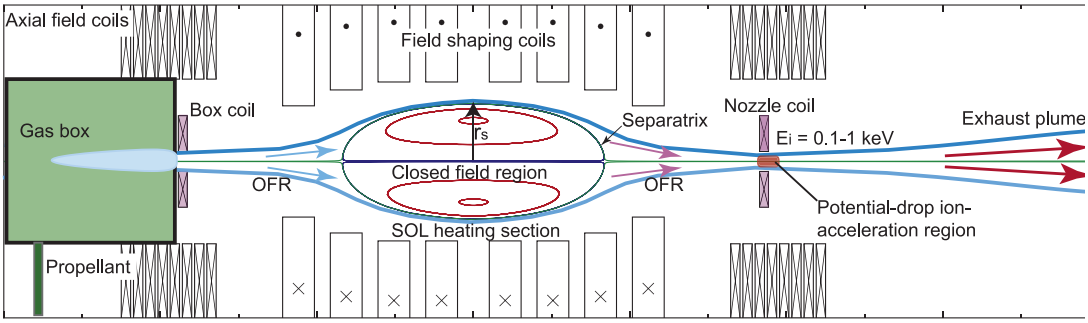


Fig. 8.9 The direct fusion drive. Deuterium from the gas box is heated as it flows around the closed field region in which the fusion occurs. OFR = open field region; SOL = scrape-off layer. Credit: Used with permission of Elsevier Science & Technology Journals, from Razin et al. (2014), conveyed through Copyright Clearance Center, Inc.

The propellant, which is also deuterium, is injected into the magnetic field from the gas box where it is ionized (figure 8.9). This initially cold plasma (comprising deuteron ions and electrons) remains outside of the closed field region where the fusion is taking place. The plasma is heated by collisions with the high-energy fusion products as it flows along the axis and around the closed field region. However, as there is no physical separation, the propellant picks up some of the main fusion products (as well as products of side reactions, like tritium) from the surface of the closed field region, which is known as the 'scrape-off layer'. This now high-energy plasma propellant is directed out of the engine via a magnetic nozzle (section 8.4). More propellant can be added at this point to increase the thrust at the expense of a lower effective exhaust velocity (a technique we will look at in section 8.8.3).

The DFD is in principle scalable to various powers of 1–100 MW. Suggested uses range from Earth-orbit manoeuvres through solar system missions to interstellar travel. Cohen et al. (2019) outline a mission to Alpha Cen using a 100 MW engine. When configured to produce an effective exhaust velocity of 15 000 km s^{-1} (0.05 c), this would ideally produce a thrust of $F = 2P/v_e = 13.3$ N (equation 3.23), although the authors assume a power-to-thrust efficiency of just 0.3, so this is lowered to 4 N. If a specific power of 25 kW kg^{-1} can be achieved, they compute the spacecraft would take about 500 years to get to Alpha Cen. (The specific power is defined here as the power of the engine divided by the mass of the engine, thus excluding the mass of the fuel, propellant, and payload.) However, this specific power is far above the published targeted specific powers of the DFD, which are around 0.3–1.5 kW kg^{-1}. A current design with a plasma volume of 2.2 m^3 and total mass of 14.5 t is predicted to produce 14.5 MW fusion power, of which 10 MW could be utilized, giving 0.7 kW kg^{-1}.

The drawback of reactors currently being developed for Earth-based power generation is that they require massive structures – magnets or lasers – that are impractical for a rocket. Conventional magnetic confinement as done in a tokamak aims to hold a relatively low density plasma at high temperatures for a long period of time. Inertial confinement, in contrast, aims for extremely rapid compression to high densities to achieve fusion far from equilibrium. Some concepts for achieving sustained fusion reactions in a rocket seek to combine aspects of both confinement approaches, an approach known as *magnetized target fusion*.

A rocket concept that uses this combined approach is the *helicity drive*. This uses a relatively modest magnetic field strength (0.3–1 T) to confine a plasma at an intermediate density

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

for an intermediate length of time. While in this state, the plasma is compressed in pulses. This inertial confinement part is done using a specially shaped magnetic field, rather than lasers. As the confinement takes place at higher densities, the engine should be more compact than a tokamak (or similar device), and as the compression can be done over a longer timescale, a less powerful pulsing system is needed. The whole system is cyclic, with pulses applied at a rate of about 1 Hz.

A key idea of the helicity drive is to use magnetic reconnection to heat the ions more efficiently than conventional approaches, by not transferring as much energy to the electrons and so reducing losses due to bremsstrahlung. One consequence of the approach is that the fusion triple product scales with the number of plasma sources N as $N^{3/2}$. This means that fusion can be achieved by simply delivering more plasma. Different fuels and fusion reactions are possible, but D+^{3}He and D+D are the favoured choices.

There have been several other investigations into fusion engines. The *fusion drive rocket* uses a convergent array of magnetic fields to compress a plasma to ignite fusion. This directly heats a solid lithium propellant. A solid propellant is favoured over a liquid or gas propellant because it saves mass on tanks. The goal is to achieve an effective exhaust velocity of order 30 km s^{-1} in a 100 t spacecraft, making this an interesting application for interplanetary missions, for example travelling to Mars in 30 to 90 days.

Whether or when any of these ideas will see development even on modest scales for space missions is, as of mid 2025, an open question. Some startups, as well as more established companies and governments, are investing into fusion power as part of a long-term move away from fossil fuels. Some of these endeavours are more serious than others, and many will no doubt produce neither fusion power nor rockets. Yet as a long-term, sustainable energy source on the Earth, as well as a high-power source for interstellar travel, fusion power is one of the relatively few options that human civilization will eventually have to develop.

8.8 Antimatter propulsion

The fission of uranium-235 releases about 0.1% of its rest energy into the kinetic energy of its products. The D+^{3}He fusion reaction releases about 0.4%. Matter–antimatter annihilation can convert a much larger fraction of the mass into potentially useable energy.

8.8.1 Main reactions

For almost every type of particle there exists an antiparticle which has many of the same properties as the particle, such as mass, but also some opposite quantities, such as charge. A collision between a particle and antiparticle results in their mutual annihilation, with their masses released as photons or as particles with varying degrees of kinetic energy.

The antiparticle of the electron is the positron, the first antimatter particle to be predicted and discovered. When an electron and position collide at low energies, they annihilate to produce two or more photons:

$$\mathrm{e}^+ + \mathrm{e}^- \rightarrow 2\gamma \ (+ \text{ more } \gamma). \tag{8.11}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

The rest energy of an electron is 0.511 MeV, so if two photons are produced (by far the most likely outcome) then each will have a wavelength of $hc/E = 0.0024$ nm. Any initial kinetic energy that the particles have will result in shorter photon wavelengths. The two photons will be moving in opposite directions, so to use them as a means of propulsion in a photon rocket (section 7.7), they need to be collimated into a narrow beam by means of reflection, which is not easy at these very short wavelengths.

If the electron and positron collide at very large energies – in excess of 105 MeV – then other particles such as muons and antimuons can be produced. At even higher energies other exotic particles including bosons are produced, as is done in large particle colliders for research purposes. But such energies would not be available aboard a spacecraft, so such results can be ignored for our purposes.

The primary matter–antimatter reaction of interest for propulsion is the annihilation of protons by antiprotons. This is more complex than the above reaction because the proton is made up of three quarks (and the antiproton of three antiquarks). At low energies the most probable reaction is

$$\mathrm{p}^+ + \mathrm{p}^- \rightarrow n\pi^+ + n\pi^- + m\pi^0 \quad (2n + m \simeq 5). \tag{8.12}$$

This produces several pions (pi-mesons), which are hadrons comprising two quarks. Due to charge conservation, the same number of positive and negative pions must be produced, and there is often also a neutral pion. The actual number varies, but in what follows we will assume that on average there are two neutral pions and 1.5 of each of the charged pions ($m = 2, n = 1.5$).

Pions are unstable. The mean lifetimes (in their rest frames) are $\tau = 2.6 \times 10^{-8}$ s for the charged pions and $\tau = 10^{-16}$ s for the neutral pion. These decay as follows:[8]

$$\pi^0 \rightarrow 2\gamma, \tag{8.13a}$$

$$\pi^+ \rightarrow \mu^+ + \nu_\mu, \tag{8.13b}$$

$$\pi^- \rightarrow \mu^- + \overline{\nu}_\mu \tag{8.13c}$$

where μ^+ and μ^- denote the muon and antimuon (which are leptons), and ν_μ and $\overline{\nu}_\mu$ denote the (muon) neutrino and antineutrino. These are all fundamental particles, but whereas the neutrinos and the photons are stable, the muons are not. They have a mean lifetime of $\tau = 2.2\,\mu$s and decay as

$$\mu^+ \rightarrow \mathrm{e}^+ + \nu_\mathrm{e} + \overline{\nu}_\mu, \tag{8.14a}$$

$$\mu^- \rightarrow \mathrm{e}^- + \overline{\nu}_\mathrm{e} + \nu_\mu\,. \tag{8.14b}$$

Once all the decays have occurred the overall reaction of the proton and antiproton is

$$\mathrm{p}^+ + \mathrm{p}^- \rightarrow n(\mathrm{e}^+ + \mathrm{e}^-) + 2m\gamma + \text{ neutrinos}. \tag{8.15}$$

8.8.2 Use in a rocket

With the proton–antiproton reaction we have a few options of which products to use. Considering the initial reaction in equation 8.12, the neutral pion is hard to couple to the spacecraft. As we shall see below, this receives about a third of the kinetic energy. We could in principle

[8] The charged pions can also decay to positron/electron and electron neutrino/antineutrino, but this is much rarer.

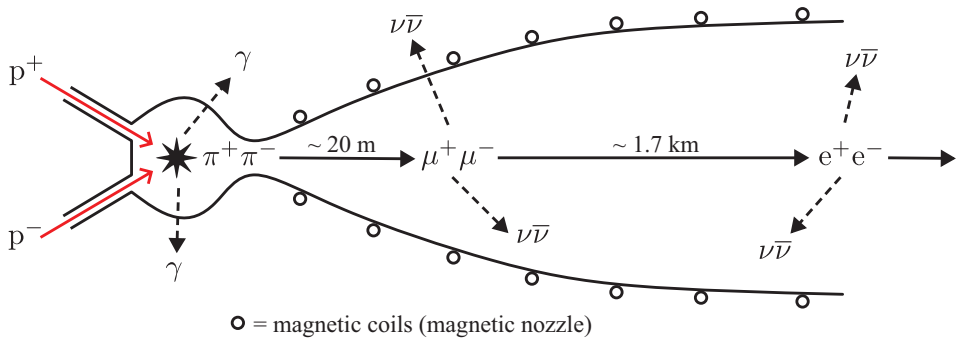
Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

make use of the photons it decays into (equation 8.13), but as these have energies of the order of hundreds of megaelectronvolts, and thus wavelengths of the order of 10^{-5} nm, they are very hard to direct. Thus about a third of the total energy from the initial proton–antiproton reaction is lost, and we focus instead on capturing the energy of the charged particles. We can either make use of the charged pions before they decay, or of their resulting muons after they decay. The former is preferable because after the decay, some of the pions' energy is given to the neutrinos, which barely interact with matter and so their kinetic energy would be lost.

For low-energy interactions, the total energy available from the initial proton–antiproton pair is the sum of their masses, which is $2 \times 938.3\,\text{MeV} = 1876.6\,\text{MeV}$. As the number of charged pions produced in the first reaction varies, so do their energies, so they have a broad spectrum of total energies. The mean is $E = 389.1\,\text{MeV}$, of which $E_0 = 139.6\,\text{MeV}$ is the rest energy. Assuming three charged pions ($2n = 3$ in equation 8.12) are produced with these energies, then 40% of the rest energy of the fuel goes into the kinetic energy of these pions. As $E = \gamma E_0$, the pions move with a speed of $\gamma = 2.8$ or $\beta = 0.93\,c$ on average. This is their speed relative to the centre-of-mass of the reactants, which is the spacecraft's rest frame. The distance they travel depends on their lifetime in the spacecraft's rest frame, which is $t = \gamma\tau$ from the Lorentz transformations (section 7.2). Thus the distance travelled in the spacecraft rest frame is

$$L = vt = \beta c\gamma\tau = c\tau\sqrt{\gamma^2 - 1} = c\tau\sqrt{\frac{E^2}{E_0^2} - 1} \tag{8.16}$$

using the relation for $\gamma\beta$ from equation 7.4. We find that charged pions with the mean energy from the reaction travel 20.3 m before decaying. To gain momentum from them, the spacecraft must therefore couple to them within this distance. This would typically be done via a magnetic nozzle as shown in figure 8.10. From equation 8.2 we can estimate the field strength required to collimate them. Charged pions have the same charge as the proton and a mass of $0.15\,m_u$. Thus to turn these pions in a radius of 1 m requires a field strength of $B = 1.2\,\text{T}$. Because the pion energy distribution has a long tail to large energies, most of the pions have smaller energies than the mean value, and so don't move as far before they decay. So they need to be coupled over less than 20 m, although their lower energies also mean their turning radii are smaller.



An antimatter engine based on proton–antiproton annihilation to pions. The momentum of the pions can be coupled to the spacecraft by magnetic fields, which collimate the pions to create a magnetic nozzle. After about 20 m the pions decay to muons, which after another 1.7 km decay to electrons. Adapted from Morgan (1982) and Forward (1982). **Fig. 8.10**

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

Table 8.2 Rest and kinetic energies for the hydrogen–antihydrogen antimatter reaction, given by equation 8.12 for the nucleons and equation 8.11 for the electron and positron. Data taken from Table 2 of Frisbee (2003a).

Particle	Rest energy [MeV]	Kinetic energy [MeV]
Fuel		
p^+	938.3	0
p^-	938.3	0
e^+	0.5	0
e^-	0.5	0
Reaction products		
$2\,\pi^0$	269.9	439.1
$1.5\,\pi^+$	209.4	374.3
$1.5\,\pi^-$	209.4	374.3
$2\,\gamma$	0	1.0

If we wait until the charged pions decay via equation 8.13, the resulting muons again have a range of energies, with an average of 298.0 MeV of which 105.7 MeV is the rest energy. The average neutrino energy is therefore $389.1 - 298.0 = 91.1$ MeV. Thus about 37% of the kinetic energy of the charged pion will be lost if we don't couple it before it decays. About 31% of the rest energy of the original proton–antiproton fuel ends up as the kinetic energy of the muons. The muons have $\gamma = 2.8$, $\beta = 0.93$, more or less the same as the pions, but as their lifetimes are longer the muons travel an average of 1.7 km before decaying.

If we wait even longer, the muons will decay to electrons, positrons, and neutrinos via equation 8.14. The neutrinos carry off about two-thirds of the energy. Many of these will annihilate via equation 8.11, leaving photons that we could then try to harness as propellant in a photon rocket.

Whether we use the pions or muons to propel the rocket, in both cases they are moving at about $0.93\,c$ relative to the rocket on average. We learned in section 7.8 that when a non-negligible amount of fuel mass is converted into energy, the effective exhaust velocity is not equal to the particle velocity. In the present case a significant fraction of the initial mass goes into neither propellant mass nor propellant energy. To be able to apply the relativistic rocket equation (equation 7.27) when we have losses, we need to compute the appropriate effective exhaust velocity.

Table 8.2 shows the typical energies of the particles involved in the reaction of hydrogen with antihydrogen to produce pions, adopting $n = 1.5$ and $m = 2$ in equation 8.12. Let us assume that magnetic fields are used to couple all of the momentum of the two charged pions to the spacecraft, but that the neutral pion is unusable. The initial energy is the sum of the initial rest and kinetic energies in the fuel, which is $2 \times (938.3 + 0.5) = 1877.6$ MeV (we now include the electrons and positrons too). Referring back to figure 7.9, a fraction $1 - \epsilon$ of this energy goes into the rest energies of the propellant particles, the charged pions. Thus

$$1 - \epsilon = \frac{209.4 + 209.4}{1877.6} = 0.2231 \quad \Rightarrow \quad \epsilon = 0.7769. \tag{8.17}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

The fraction of the initial energy that goes into propelling these particles (and potentially any photons) is $\eta\epsilon$. This is equal to the kinetic energy of the charged pions, assuming that the gamma photons are not used to propel the spacecraft (so they will be counted as part of the losses). Thus

$$\eta\epsilon = \frac{374.3 + 374.3}{1877.6} = 0.3987, \tag{8.18}$$

from which it follows that $\eta = 0.5132$. We are assuming $\zeta = 1$ (no photons used for propulsion), so putting these three numbers into equation 7.39 we get $v_e = 0.5804\,c$.

The effective exhaust velocity is less than the velocity of the pions, calculated above to be $0.93\,c$. This is because of the energy lost to the neutral pions (and a tiny amount to the photons). These losses are a fraction $(1 - \eta)\epsilon = 0.3782$ of the initial energy, which is 710.1MeV. This is the sum of the rest and kinetic energies of the neutral pions and the energy of the photons. If we could somehow couple also the neutral pions to the spacecraft, then we would achieve

$$\epsilon = 1 - \frac{209.4 + 209.4 + 269.9}{1877.6} = 0.6332, \tag{8.19}$$

and no losses (we ignore the small energy of the gamma), so $\eta = 1$, giving $v_e = 0.93$, which is indeed the velocity of the pions.

8.8.3 Thrust augmentation

We have found that exploiting the charged pions from hydrogen–antihydrogen reactions can theoretically deliver $v_e = 0.58\,c$. Using this to propel a rocket to a velocity of $0.1\,c$ so we can get to Alpha Cen within 50 years would need an initial-to-final mass ratio of only 1.19 from the relativistic rocket equation (the classical one gives virtually the same result). Thus just 16% of the total rocket mass needs to be fuel.

However, the dry spacecraft mass (everything minus fuel/propellant) will not be small on account of the engine (fuel containment, magnetic nozzle, etc.). So if the dry spacecraft mass were 100 t (surely optimistic), we would still need 19 t of fuel, which is 9.5 t of antihydrogen. This is 15 orders of magnitude more than the current annual production (see section 8.8.4).

Assuming that the amount of antihydrogen available is limited, we can achieve a larger thrust by mixing in more hydrogen and letting the high-energy pions heat the excess non-annihilated hydrogen through collisions. This will of course lower the effective exhaust velocity. But it turns out this can still produce a larger Δv overall.

This idea is quite general and not specific to antimatter or even nuclear reactions. To calculate the resulting Δv, we turn again to the general relativistic formulation for the mixed matter–photon rocket in section 7.8, now for the case of a matter rocket ($\zeta = 1$). Suppose the total mass of fuel is m_a. For the antimatter rocket this is half hydrogen and half antihydrogen. Let m_p be the mass of inert propellant that will be added to the exhaust; it does not undergo any reactions and does not release energy. The fraction of the initial mass that is ejected as propellant is

$$\frac{(1 - \epsilon)m_a + m_p}{m_a + m_p} = \frac{1 - \epsilon + x}{1 + x}, \tag{8.20}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

where $x = m_p/m_a$. Using this in place of $(1 - \epsilon)$ in the conservation of momentum (equation 7.36), the effective exhaust velocity (equation 7.37b) becomes

$$v_e = \left(\frac{1 - \epsilon + x}{1 + x}\right) \gamma_m v_m. \tag{8.21}$$

In the conservation of energy equation (equation 7.38b), the mass fraction to be accelerated is $1 - \epsilon + x$ with the same amount of energy as before, and so the velocity of the particles is given by

$$\gamma_m = \frac{\eta\epsilon}{1 - \epsilon + x} + 1. \tag{8.22}$$

Continuing as before, we find that the effective exhaust velocity of the thrust augmented matter rocket is

$$\frac{v_e}{c} = \frac{\sqrt{\eta\epsilon(\eta\epsilon + 2[1 - \epsilon + x])}}{1 + x}. \tag{8.23}$$

When $x = 0$ we recover the non-augmented result. When $\epsilon \ll 1$ or $x \gg 1$ we get

$$\frac{v_e}{c} \simeq \sqrt{\frac{2\eta\epsilon}{1 + x}} \quad (\epsilon \ll 1 \text{ or } x \gg 1). \tag{8.24}$$

This is the same result that we get for the velocity of propellant particles when we do a classical calculation in which a fraction $\eta\epsilon$ of the rest energy of mass m_a provides the kinetic energy of the remaining mass, that is,

$$\eta\epsilon m_a c^2 = \frac{1}{2}([1 - \epsilon]m_a + m_p)v_e^2, \tag{8.25}$$

which, being classical, requires $\epsilon \ll 1$.

Let us use these results to compare how the Δv attained by a thrust augmented rocket compares with that of a non-augmented rocket. The mass ratio is $R = (m_d + m_a + m_p)/m_d$, where m_d is the dry mass of the rocket (without fuel or propellant), and for the non-augmented rocket $m_p = 0$. Using this and the appropriate values of v_e in the relativistic rocket equation (equation 7.27), we can compute Δv. Figure 8.11 shows the case for $m_d = 100\,\text{t}$ and $m_a = 1\,\text{t}$ (although we can use any units as only the mass ratios are relevant) for increasing R, that is, increasing m_p. We see that although augmenting the thrust decreases v_e, it significantly increases Δv. The non-augmented rocket has $v_e = 0.5804\,c$, $R = 1.01$, and $\Delta v = 0.0058\,c$. If we augment this with $m_p = 1\,\text{t}$ of inert propellant ($R = 1.02$), v_e only decreases slightly to $0.5325\,c$, but Δv nearly doubles to $0.0105\,c$. If we augment with $m_p = 393\,\text{t}$ ($R = 4.94$), we achieve $\Delta v = 0.0718\,c$ at $v_e = 0.0449\,c$. This is the optimal amount of inert propellant to add (the maximum of the curve). Adding more inert propellant then lowers Δv, because the increase in ejected mass is now being offset faster by the decrease in v_e.

We see that adding inert propellant increases Δv compared to just using the antimatter. In practice we need more mass for the propellant tanks and pumps, which would increase m_d, but a net increase in Δv should still be achievable. Adding inert propellant won't give a larger Δv than when burning that additional propellant as fuel, of course. But adding inert propellant becomes a useful option if the fuel is hard to acquire (as in the case of antimatter) or if we cannot react it fast enough to get a large thrust (as might be the case with fusion).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

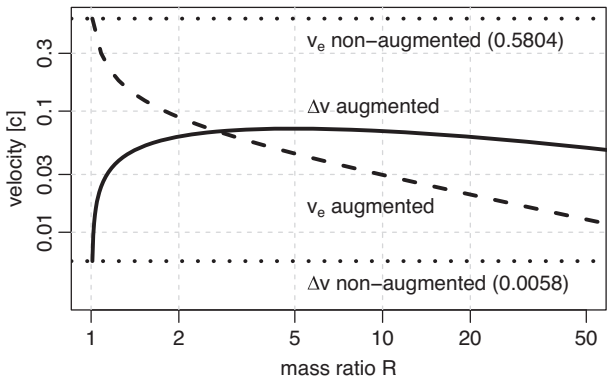


Fig. 8.11

The thrust augmented hydrogen-antihydrogen rocket driven by the charged pions (equation 8.12, which gives $\epsilon = 0.7769$, $\eta = 0.5132$, $\zeta = 1$) with a dry mass of $m_d = 100$ and unit mass matter+antimatter fuel ($m_a = 1$; the units are arbitrary). The dashed line shows how the effective exhaust velocity (equation 8.23) varies as we increase the mass m_p of inert propellant added to the reactant, in terms of the rocket's mass ratio $R = (m_d + m_a + m_p)/m_d$. The solid line shows the corresponding Δv achieved by the rocket according to the relativistic rocket equation. The two horizontal dotted lines show the same quantities for the non-augmented rocket, where $m_p = 0$, which here are independent of R. Note the log scales on both axes.

8.8.4 Practicalities

Use of antimatter as an energy source presents two significant challenges. The first is that it hardly exists in the local universe. Some is produced via interactions with cosmic rays, and positrons are produced in some beta decay processes. Antiprotons can be produced in particle accelerators by firing protons at very high energies into a dense target, such as tungsten. The resulting antiprotons have a very large spread in energies, so the beam must be cooled to be able to store them in a magnetic ring. Current production rates are extremely small. Accelerators produce of order 10^{16} antiprotons (10^{-11} kg) per year, although of course the experiments are not designed to produce them on industrial scales. There are ideas to increase annual worldwide production to around 0.02 kg, but this is still far below the amount needed to propel a fast spacecraft of any size.

The second problem with antimatter is its storage. As a spacecraft would be made from normal matter, storage cannot involve contact and would have to use electric or magnetic fields instead. Antiprotons are currently stored in motion in circular storage rings, but only for short periods (accelerated charge loses energy through synchrotron radiation), and these magnetic systems are likely to be too massive for a spacecraft. One idea for high-density storage is as molecular antihydrogen ice suspended in a magnetic field, which must be held at very low temperatures (1–2 K) to prevent evaporation. Small amounts could be extracted and channelled electromagnetically to the combustion chamber as required.

One advantage of antimatter propulsion is that the reaction occurs spontaneously. Unlike fusion reactors, we do not need massive equipment for inertial or magnetic confinement to initiate or sustain the reaction. This mass saving could offset some of the mass needed for the antimatter confinement.

Antimatter has an alternative potential use in propulsion as a means of triggering nuclear fission in subcritical masses. As an antiproton has a negative charge, it can penetrate the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

nucleus of a uranium atom. Here the antiproton annihilates with a proton (equation 8.12) and the resulting high energy pions can induce nuclear fission. This can even induce fission in uranium-238. The resulting energy and momentum can then propel a rocket in the same way as either a fission fragment rocket (section 8.3.2) or a fission pulse rocket (section 8.3.3). The advantage of the antimatter-induced fission, however, is that it no longer needs a critical mass, and so fission can be achieved with very small amounts of fuel. Only microgrammes of antihydrogen may be necessary to induce the fission. This approach is particularly attractive for the deceleration of spacecraft at its target destination.

8.9 Summary

Nuclear reactions release about 10^7 times more energy per unit mass than chemical reactions, and so in principle can deliver in the order of $10^{3.5} = 3000$ times larger effective exhaust velocities (as $v \sim \sqrt{K}$).

The products of nuclear reactions could be exhausted directly from the rocket. As the products have extremely high energies, they cannot be collimated with a conventional nozzle. The charged products can instead be coupled to the spacecraft using a magnetic nozzle.

One example of direct use is the fission fragment rocket, which can theoretically achieve $v_e = 0.039\,c$ with uranium-235 fuel. The thrusts tend to be low, however. The fusion pulse rocket instead extracts energy and momentum from a rapid sequence of small fusion explosions. Using the D+^{3}He reaction, this can theoretically deliver $v_e = 0.089\,c$. The Daedalus and Icarus projects have outlined concepts that could get to Alpha Cen in 50–100 years. A big drawback of this reaction, though, is the lack of helium-3 on Earth, which would have to be mined from the Moon or Jupiter.

The energy from a nuclear reaction could instead be used to heat an inert propellant, as in the thermal fission rocket or direct fusion drive. In the former case this lowers the effective exhaust velocity significantly, but can provide quite large thrusts.

Most nuclear reactions put a lot of energy into neutrons or produce bremsstrahlung, both of which will heat and damage the spacecraft. Shielding and massive large radiators for cooling will be essential.

Antimatter reactions release much more energy than fission or fusion. The annihilation of protons with antiprotons to produce pions, for example, can provide a v_e of up to $0.580\,c$. However, only tiny amounts of antimatter can currently be produced, so this remains a very far-fetched concept.

Perhaps counter-intuitively, the Δv of a rocket can be increased by mixing in an inert propellant, such as hydrogen, a process known as thrust augmentation. Although this lowers the effective exhaust velocity, it can increase the mass flow rate sufficiently to more than compensate for this and the extra propellant mass.

Nuclear fusion is one of the best options we have for accelerating spacecraft to 3–9% the speed of light, and so to get to Alpha Cen in 50–100 years. Such spacecraft would necessarily be massive (on the order of thousands of tonnes, or more), much of which would be fuel and propellant and probably cooling systems, but they could then carry fairly substantial payloads.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009

8.10 Exercises

1. Following the approach in section 8.3.2, work out the initial mass and propellant power of a fission fragment rocket that can get to Alpha Cen within 100 years. Assume the engine has an effective exhaust velocity $v_e = 0.017\,c$ and that the total initial mass is 200 t plus the mass of the propellant. Write code to compute the time taken, using the relevant equations from chapter 3. There is no unique answer – the spacecraft could accelerate rapidly then spend most of the journey cruising, or it could accelerate more leisurely for most of the journey – so it's probably easiest just to use trial and error. This exercise should give a feel for the trade-offs involves, and an appreciation for the magnitudes of the quantities.
2. Using both approaches described in section 8.5.2, compute the effective exhaust velocity for the catalyzed deuterium reaction of equation 8.9 assuming the neutron (a) can, and (b) cannot, be coupled to the rocket.
3. In section 8.8.3 we looked at thrust augmentation in an antimatter engine. The concept applies to a classical engine too, for example mixing inert cold hydrogen into a hot exhaust. Suppose a non-augmented engine exhausts mass m_a at rate $\dot{m}_a$ and velocity v_a. We now mix in an inert propellant of mass m_p at rate $\dot{m}_p$. (i) Show that the effective exhaust velocity becomes $v_a/\sqrt{1 + m_p/m_a}$. (ii) Show that this is the same as equation 8.24 when $\eta = 1$. (iii) Compute the thrust for the non-augmented and the augmented cases. Show that the thrust in the augmented case is larger when $\dot{m}_p/\dot{m}_a > \sqrt{1 + m_p/m_a} - 1$. (iv) Equation 5.10 predicts that adding a low-mass inert propellant – in order to lower the molar mass of the exhaust – *increases* the effective exhaust velocity. How can we reconcile this with the answer to part (i)?

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 12 Nov 2026 at 20:38:10, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.009