Laser sails 11

The velocity that can be attained by a solar sail is limited by the Sun's intensity. In this chapter we examine replacing the Sun with a much more intense artificial radiation source, such as a laser, to achieve a higher velocity. We look more closely at the mechanics of photons incident on sails, taking into account relativistic effects and the finite travel time of light. Diffraction of the directed radiation source is a major issue, because it limits the amount of light intercepted by the sail, requires large optics, or demands high laser power, or some combination of these. We look into achieving the kilometre-scale optical systems required using aperture synthesis and adaptive optics. Given the large laser powers involved (tens to hundreds of gigawatts), sail heating becomes a limiting factor, and the choice of materials is critical. Having looked at the theoretical issues, I summarize a few of the numerous laser sail designs that have been published that bring together many of the issues discussed.

11.1 Relativistic kinematics of a laser sail

We start with the fundamental kinematics of a photon-propelled sail. Analyzing the momentum and energy transfer at the sail, initially for perfect reflection, we derive an expression for the velocity of the sail as a function of the laser energy. We then derive the equation of motion of the sail allowing for imperfect reflection, and solve this for a constant power laser. In this section we neglect diffraction loses.

11.1.1 Reflection of photons off a relativistic sail

Consider a sail receding directly away from a laser at velocity $\beta > 0$, as shown in figure 11.1. The laser emits a single photon of frequency ν_e in the laser rest frame. In the sail rest frame, the laser appears to be receding, so when received, this photon is Doppler shifted to a lower frequency of $\nu'_e = \nu_e\sqrt{(1-\beta)/(1+\beta)}$ (see section 9.2).[1] If we assume the photon is perfectly reflected by the sail, then in the sail rest frame the reflected photon still has a frequency of ν'_e. But when viewed from the laser rest frame, which sees a receding sail, this photon is Doppler shifted again, and by the same factor if the sail speed is unchanged. Thus the frequency of the reflected photon in the laser rest frame is

$$\nu_r = \nu_e\left(\frac{1-\beta}{1+\beta}\right). \tag{11.1}$$

The reflected photon has a lower energy. This means it must have given some of its energy to the sail. However, we assumed that the reflected photon had the same frequency as the

[1] The primed and unprimed labels are swapped with respect to the notation in chapter 9, because now the photons are emitted from the laser rest frame S.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

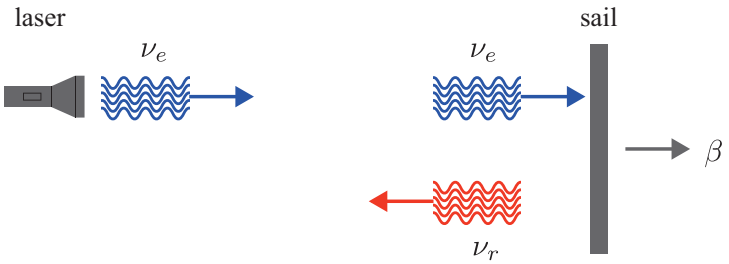


Fig. 11.1 The view from the laser rest frame of the reflection of photons off a sail receding at velocity β.

incident photon (in the sail rest frame), and so did not lose energy. Clearly this assumption is incompatible with the assumption that the speed (and thus kinetic energy) of the sail did not change. This is similar to the classical problem of a ball bouncing elastically off a wall, in which the ball returns with the same speed. This can only happen if the wall is infinitely massive, otherwise the ball will give energy to the wall, and so return more slowly.

We will deal with this dilemma correctly below, but in practice the sail has a rest mass–energy that is much larger than the energy of a single photon. In that case, equation 11.1 is a very good approximation. It tells us that the larger the sail velocity, the larger the redshift, and so the larger the fraction of energy that has been transferred to the sail. In other words, the efficiency of energy transfer from laser to sail increases with increasing velocity.

We might stop at this point and ask how energy conservation works with the Doppler effect, before we even get to the problem of reflection. Specifically, if the emitted photon has energy $h\nu_e$ in the laser rest frame, but energy $h\nu_e'$ in the sail rest frame, where has the energy gone? Energy, like velocity and momentum, is a frame-dependent quantity, so it is not conserved between different rest frames. This is true in classical systems also: An object at rest has zero kinetic energy, but it has a non-zero kinetic energy in a reference frame that it moving relative to it. The velocity of a photon is invariant across all reference frames, but its energy (and momentum) is not.

Let us now consider more rigorously the problem of a single photon reflecting off a sail, from the viewpoint of the laser rest frame. The sail has mass m and initial velocity β_i. A photon of frequency ν_e approaches the sail, then is reflected in the opposite direction with frequency ν_r. This exchange of momentum and energy results in the sail changing its velocity to β_f (its mass is assumed unchanged). Recall that for photons $E = h\nu = pc$, where p is the momentum. Equating the initial and final momentum and energy, we get two conservation equations in two unknowns, ν_r and β_f:

$$\frac{h}{c}\nu_e + \gamma_i mc\beta_i = -\frac{h}{c}\nu_r + \gamma_f mc\beta_f \quad \text{(momentum)}, \tag{11.2a}$$

$$h\nu_e + \gamma_i mc^2 = h\nu_r + \gamma_f mc^2 \quad \text{(energy)}. \tag{11.2b}$$

We rearrange each of these equations into something more convenient:

$$\frac{h}{mc^2}(\nu_e + \nu_r) = \gamma_f\beta_f - \gamma_i\beta_i \quad \text{(momentum)}, \tag{11.3a}$$

$$\frac{h}{mc^2}(\nu_e - \nu_r) = \gamma_f - \gamma_i \quad \text{(energy)}, \tag{11.3b}$$

and then sum them to eliminate ν_r:

$$\frac{2h\nu_e}{mc^2} = \gamma_f(1+\beta_f) - \gamma_i(1+\beta_i). \tag{11.4}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

Using the identity in equation 7.6 and defining $q = h\nu_e/mc^2$, the ratio of the photon energy to the sail rest mass–energy, this becomes

$$2q = \sqrt{\frac{1+\beta_f}{1-\beta_f}} - \sqrt{\frac{1+\beta_i}{1-\beta_i}}. \tag{11.5}$$

Rearranging this we arrive at an expression for the final velocity of the sail:

$$\beta_f = \frac{(2q+\phi_i)^2 - 1}{(2q+\phi_i)^2 + 1}, \quad \text{where} \quad \phi_i = \sqrt{\frac{1+\beta_i}{1-\beta_i}}. \tag{11.6}$$

When the sail starts at rest, $\phi_i = 1$, and we get (with $E_l = h\nu_e$)

$$\beta_f = \frac{(2q+1)^2 - 1}{(2q+1)^2 + 1}, \quad \text{where} \quad q = \frac{E_l}{mc^2}. \tag{11.7}$$

When $q = 1$, where the photon energy supplied equals the rest energy of the sail, the sail attains $\beta_f = 4/5$. For $q \ll 1$, equation 11.7 becomes $\beta_f \simeq 2q$. A sail of mass 10 g at rest reflecting a 1 μm photon has $q = h/\lambda mc = 2.2 \times 10^{-34}$ and so would acquire a velocity of order $10^{-34}\,c$. Clearly, the approximation we used to arrive at equation 11.1 earlier is valid in all practical applications. In the derivation of equation 11.7, we have assumed perfect reflection in the sense that no photons are absorbed by or transmitted through the sail, but we have *not* assumed that the photons are reflected with the same energy. Note that equation 11.7 is also the equation we would derive in the initial rest frame of the sail, with q then defined as $h\nu'_e/mc^2$ (see exercises).

The above derivation also applies if, instead of a single photon hitting the sail, we had a set of N photons. In this case $E_l = Nh\nu_e$ is the total energy (in the laser rest frame) emitted by the laser that arrives at the sail. It turns out, as will be shown below, that the equation for β_f still applies if the photons arrive at different times when the sail has a range of velocities. That is, the final velocity depends only on the total amount of energy received, and not on how it is received over time. The variation of the sail velocity with q is shown in figure 11.2. A 10 g sail (rest energy 9.0×10^{14} J) propelled by a 100 GW laser for 1000 s (10^{14} J) attains a final velocity of $0.20\,c$, as does a 100 g sail powered for 10 times as long.

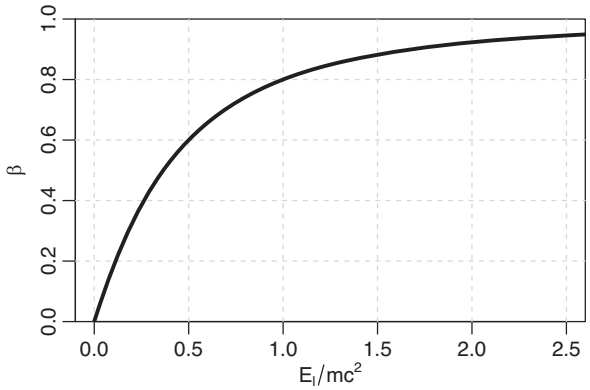


Fig. 11.2 The dependence of the velocity acquired by a sail of mass m when accelerated from rest by photons of total energy E_l (equation 11.7). The curve asymptotes to $\beta = 1$ as $q = E_l/mc^2$ increases to infinity. This assumes all photons reach the sail and are reflected (zero transmission through, or absorption by, the sail).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

11.1.2 Energy efficiencies for perfect reflection

How much of the laser energy is put into the kinetic energy of the sail? The energy provided by the laser (in the laser rest frame) and received by the sail is $E_l = qmc^2$. From equation 11.5 with $\beta_i = 0$, and using the identity in equation 7.6, we can write this in terms of the velocity as

$$E_l = \frac{1}{2}[\gamma(1+\beta) - 1]mc^2. \tag{11.8}$$

As the kinetic energy (KE) of the sail is $(\gamma - 1)mc^2$, the overall energy efficiency is the ratio of these:

$$\text{overall efficiency} = \frac{\text{KE}_{\text{sail}}}{E_l} = \frac{2(\gamma - 1)}{\gamma(1+\beta) - 1}. \tag{11.9}$$

This is plotted in figure 11.3. We see that the efficiency increases monotonically with increasing velocity. At very low velocities, the efficiency is near zero. Most of the energy is reflected off the sail. This is consistent with equation 11.1, which tells us that at low velocity the reflected photon has a very similar frequency, and thus energy, to the incident photon.[2] As the velocity increases, more energy is put into the sail, and correspondingly the reflected photon is redshifted more.

It is interesting at this point to compare the laser sail with the photon rocket, discussed in section 7.7. The energy E_r that the photon rocket has to expend to get its velocity to β is given by equation 7.35. This is twice the amount that the laser has to deliver (E_l in equation 11.8) to give the sail the same kinetic energy. Although we can in principle achieve near unit energy efficiency with a laser sail, with a photon rocket the maximum efficiency is only 1/2. This is in keeping with the disadvantage of a rocket having to carry its energy/propellant source. Observe that whereas E_l is measured in the laser rest frame, E_r was measured in the photon rocket's (accelerating, instantaneous) rest frame. But in both cases the amount of

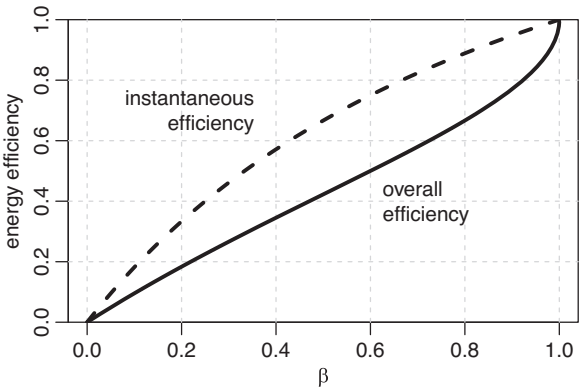


Fig. 11.3 The variation of the energy transfer efficiencies of a laser sail as a function of the sail's velocity. The solid line shows the overall efficiency, equation 11.9, and the dashed line the instantaneous efficiency, equation 11.10 .

[2] Recall that β is the velocity of the sail after the first photon has hit it, so the energy transfer and efficiency are not exactly zero for the first photon.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

energy expended to achieve the motion is measured in the rest frame in which the energy is expended.

This factor of two difference also occurs when we consider the momentum. The photon rocket is propelled by photons emitted from it, whereas the laser sail is propelled by photons reflected off it, which, for perfect reflection, give twice the change in momentum as a reflected photon.

As we are primarily interested in accelerating sails to speeds of 0.1–0.2 c, figure 11.3 shows that most of the energy sent by the laser will be reflected back to the laser. When the laser power is being sent by a large telescope (discussed below), we could simply reflect the (redshifted) photons back to the spacecraft, where they can again transfer some of their kinetic energy to the spacecraft. This can in principle be repeated more than once, as long as the spacecraft has not attained its final velocity. But as each reflected beam will be significantly broadened by diffraction in practice, this 'photon recycling' may not help much.

The efficiency in equation 11.9 is the overall efficiency of energy transfer when accelerating the sail from rest to β. Another efficiency is that of a single photon transfer (or at least a small number of photons) when the sail is moving at velocity β. If the incident photon has energy E_e, and the reflected photon energy E_r (both in the laser rest frame), then the energy transmitted to the sail is $E_e - E_r$ and the 'instantaneous efficiency' is $(E_e - E_r)/E_e$. With $E = h\nu$ and using equation 11.1, we get

$$\text{instantaneous efficiency} = 1 - \frac{E_r}{E_e} = \frac{2\beta}{1+\beta}. \tag{11.10}$$

This is plotted as the dashed line in figure 11.3.

11.1.3 The equation of motion

So far we have considered how the final velocity of the sail varies with the total amount of energy it receives, but not how its velocity varies with time or distance travelled. To determine these we need to take into account that the photons require a finite amount of time to reach the sail. Consider time measured in the laser rest frame. Because of the photon delay, the sail will achieve the velocity derived above at a time later than when the last photons were emitted by the laser. This is illustrated in figure 11.4. Let t denote the time in the laser rest frame when a photon arrives at the sail. This photon left the laser at an earlier time $T = t - x/c$, where $x \geq 0$ is the distance of the sail from the laser when the photon arrives at the sail. The distance x is also measured in the laser rest frame. The laser is switched on and the first photons are emitted at time T_A. These arrive at the sail at event A, when the time is t_A. The laser is switched off at time T_B, and the final photons arrive at the sail at event B, when the time is t_B. Even though the 'coordinate' time t and the 'retarded' time T are measured in the same reference frame, the interval $t_B - t_A$ is longer than the interval $T_A - T_B$ because of the distance travelled by the sail over that duration. Moreover, because t depends on x – the distance of the sail from the laser – these intervals will vary as the sail accelerates, as we will see below. It also follows from figure 11.4 that in the laser rest frame, the rate of photon arrival at the sail is lower than the rate of emission from the laser, because each successive photon takes longer to arrive. This is a purely classical effect resulting from the finite speed of light.

Following this discussion, we now recognize that equation 11.7 gives the velocity of the sail at time t in the laser rest frame corresponding to the energy E_l emitted by the laser up to

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

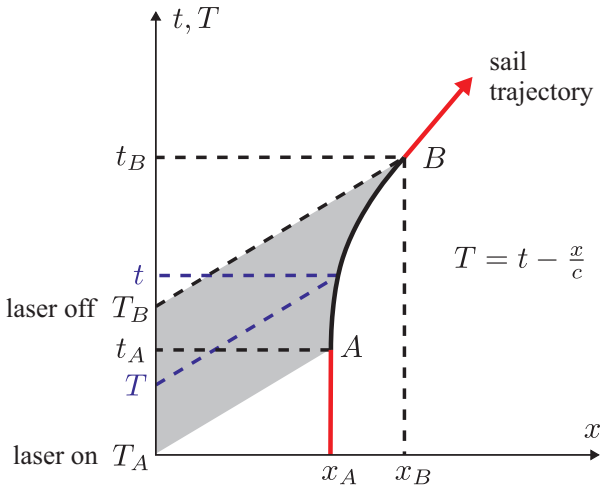


Fig. 11.4 A spacetime diagram of a laser sail, where x is the distance of the sail from the laser in the laser's rest frame. Both t and T are times in the laser rest frame. Photons emitted from the laser at time T arrive at the sail at time t, where $T = t - x/c$ is referred to as the 'retarded' time. The sail starts at $x = x_A$ and remains at rest until the first photons hit it at event A. The sail then accelerates (the curved part of its trajectory) until the last photons hit it at event B; after that the sail moves with constant velocity. The grey area shows the photons travelling from the laser to the sail.

retarded time T. If the laser continues to emit photons up until time T_B, then the final velocity of the sail – once all the photons have caught up with it – at time t_B is given by equation 11.7 with a larger value of E_l.

With this distinction between the times in mind, we now derive an equation of motion for the sail.

To evaluate the net force on the sail, we consider the reflection, transmission, absorption, and emission of photons for a relativistically moving sail. All photon momenta (equivalently frequencies, as $p = h\nu/c$) and photon rates $\dot{n}$ are defined in the laser rest frame, and all are positive.

The laser emits photons of momentum p_l at a rate $\dot{n}_l$. Because the sail is receding from the laser at a speed β, the photons arrive at the sail at a lower rate of

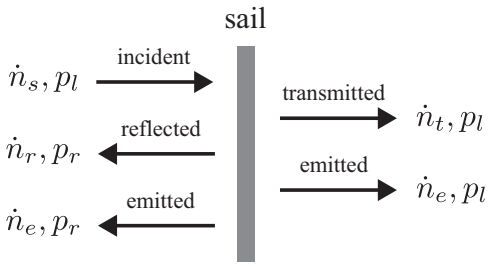
$$\dot{n}_s = \dot{n}_l(1-\beta). \tag{11.11}$$

This is a purely classical effect that we considered as part of the Doppler effect in section 9.2. The momentum (frequency) of the photons is unchanged because we are always considering the laser rest frame. These incident photons may be reflected, absorbed (and then re-emitted), or transmitted through the sail, as illustrated in figure 11.5. We now consider the effect of each of these on the sail as they enter or leave the sail. As in chapter 10, the absorptance and reflectance of the sail are denoted $\tilde{a}$ and $\tilde{r}$, respectively; the transmittance is $1 - \tilde{r} - \tilde{a}$:

- The incident photons of momentum p_l, rate $\dot{n}_s$, act to increase the momentum of the sail.
- Photons of momentum p_r are reflected at rate $\dot{n}_r = \tilde{r}\dot{n}_s$ and act to increase the momentum of the sail. These photons undergo a double Doppler shift (equation 11.1) and so

$$p_r = p_l\left(\frac{1-\beta}{1+\beta}\right). \tag{11.12}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012



A sketch of the various contributions of the photon rates and momenta to the motion of a light sail. Quantities are defined in the laser rest frame. Fig. 11.5

- Photons of (unchanged) momentum p_l are transmitted through the sail at rate $\dot{n}_t = (1 - \tilde{r} - \tilde{a})\dot{n}_s$, and so are a source of loss of momentum of the sail. (We could instead think of these as being subtracted from the incident beam.)
- Photons are absorbed at rate $\dot{n}_a = \tilde{a}\dot{n}_s$. Assuming the sail is in thermal equilibrium, these will be re-emitted at the same rate. We can consider half to be re-emitted in the forward direction as though they were transmitted (so at momentum p_l) and the other half to be re-emitted in the backward direction as though they were reflected (so at momentum p_r). In that case $\dot{n}_e = \dot{n}_a/2$.

Bringing all of the above points together, the rate of change of momentum of the sail is

$$\frac{dp}{dt} = \dot{n}_s p_l + \dot{n}_r p_r - \dot{n}_t p_l + \dot{n}_e p_r - \dot{n}_e p_l \tag{11.13a}$$

$$= \dot{n}_s p_l + \dot{n}_s \tilde{r} p_r - \dot{n}_s (1 - \tilde{r} - \tilde{a}) p_l + \dot{n}_s \frac{\tilde{a}}{2}(p_r - p_l) \tag{11.13b}$$

$$= \frac{1}{2}\dot{n}_s (2\tilde{r} + \tilde{a})(p_r + p_l). \tag{11.13c}$$

Using equations 11.11 and 11.12 we can eliminate $\dot{n}_s$ and p_r to get

$$\frac{dp}{dt} = p_l \dot{n}_l (2\tilde{r} + \tilde{a}) \left(\frac{1-\beta}{1+\beta}\right). \tag{11.14}$$

As the power of the beam in the laser rest frame is $P_l = c p_l \dot{n}_l$, we can write the expression for the force on the sail due to the laser beam as

$$F = \frac{dp}{dt} = \frac{P_l}{c}(2\tilde{r} + \tilde{a})\left(\frac{1-\beta}{1+\beta}\right). \tag{11.15}$$

Using the relativistic version of Newton's second law (section 7.4, equation 7.25) to relate this force to the coordinate acceleration $a = c\, d\beta/dt$, and neglecting any other forces, we get the equation of motion for the sail of mass m:

$$\frac{d\beta}{dt} = \frac{P_l}{mc^2}\frac{1}{\gamma^3}\left(\frac{1-\beta}{1+\beta}\right)(2\tilde{r} + \tilde{a}) = \frac{P_l}{mc^2}\frac{(1-\beta)^2}{\gamma}(2\tilde{r} + \tilde{a}). \tag{11.16}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

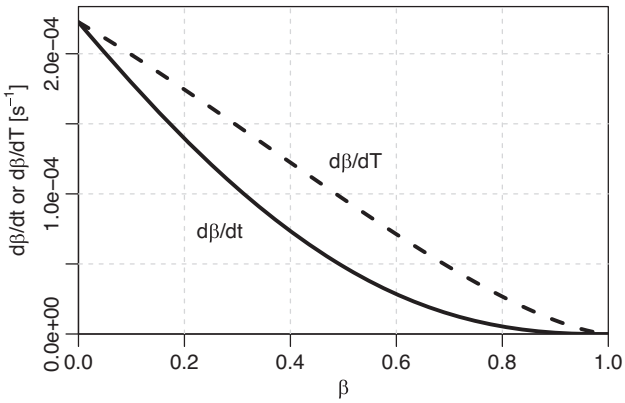


Fig. 11.6 The variation of the sail acceleration as a function of velocity ($\beta c = dx/dt$), for the acceleration defined in terms of both the coordinate time $d\beta/dt$ (solid line; equation 11.16) and the retarded time $d\beta/dT$ (dashed line; equation 11.25). This is illustrated for a 10 g sail propelled by a laser power of 100 GW, assuming perfect reflection ($\tilde{r} = 1$).

This gives the acceleration of the sail at coordinate time t when its velocity is β, where P_l is the power of the photons that have just arrived at the sail, and so were emitted from the laser at the earlier retarded time T (see figure 11.4). The variation of this acceleration with velocity is shown as the solid line in figure 11.6 for a 10 g sail accelerated with a 100 GW laser.[3] The acceleration decreases monotonically with increasing velocity. Thus even though at velocities near the speed of light the energy transfer efficiency is very high (figure 11.3), the acceleration drops to zero.

In the non-relativistic limit of $\beta \ll 1$, equation 11.16 becomes, with $v = \beta c$,

$$\frac{dv}{dt} = \frac{P_l}{mc}(2\tilde{r} + \tilde{a}) \quad \text{(non-relativistic limit)}. \tag{11.17}$$

In this case the acceleration is independent of velocity, as we would expect. This is of course the same equation that we would derive from equation 10.19 with pressure equal to $(1/A)mdv/dt$ and intensity equal to P_l/A, where A is the area of the sail.

11.1.4 The solution of the equation of motion

The integration of equation 11.16 is

$$\int \frac{1}{2\tilde{r} + \tilde{a}} \gamma^3 \left(\frac{1+\beta}{1-\beta}\right) d\beta = \int \frac{P_l}{mc^2}\, dt. \tag{11.18}$$

Let us now assume that both the sail mass and the laser power are constant, so $\int_{t_A}^{t} P_l\, dt = P_l(t - t_A)$. Assume also that the reflectance and absorptance are independent of wavelength over the range that it varies due to the increasing redshift (range of β). Integrating the left side and setting $\beta = 0$ at $t = t_A$ we get

$$t - t_A = \frac{mc^2}{3P_l} \frac{1}{2\tilde{r} + \tilde{a}} \left[\frac{(1+\beta)(2-\beta)\gamma}{(1-\beta)} - 2\right]. \tag{11.19}$$

This is the (coordinate) time that has elapsed in the laser rest frame when the spacecraft reaches velocity β.

[3] For comparison, the acceleration due to gravity from the Sun at 1 au is $d\beta/dt = 2 \times 10^{-11}\ \text{s}^{-2}$.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

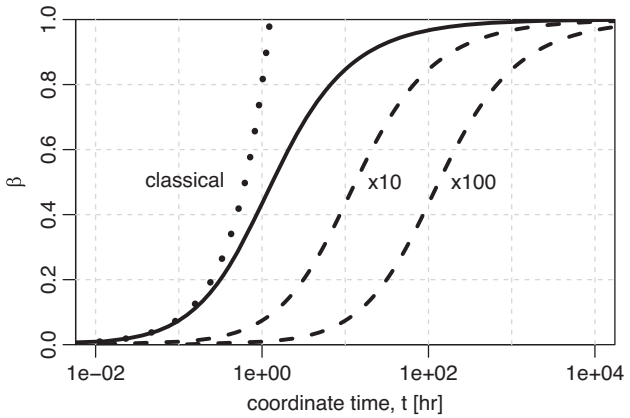


Fig. 11.7 The variation of the velocity with coordinate time t for a 10 g sail spacecraft propelled by a laser power of 100 GW assuming perfect reflection (solid line; equation 11.19). The two dashed lines show an increase in the mass (or decrease in the power) by factors of 10 and 100. The dotted line shows the classical, non-relativistic result (equation 11.20) for the nominal case.

As t is the time when a photon arrives at the sail, $P_l(t - t_A)$ is the energy that has arrived at the sail by this time. As we can see from figure 11.4, this is the energy that was emitted by the laser in the time period T_A to T. Thus $P_l(t - t_A)$ is *not* the energy that has been emitted by the laser by time t. By this time the laser has emitted more photons (those after time T), but these photons have not yet reached the sail.

Equation 11.19 does not have a closed-form solution for β in terms of t, but we can still plot the motion. This is shown in figure 11.7 for three different masses, or equivalently, three different values of the velocity-independent factor in equation 11.19. A 10 g spacecraft powered by a 100 GW laser will reach a velocity of 0.1 c after 8.3 minutes and 0.2 c after 19 minutes. Due to the linear dependence of the time on mass and power, a spacecraft 10 times more massive will take 10 times as long, for example.

In the non-relativistic limit of $\beta \ll 1$, the term in square brackets in equation 11.19 becomes 3β, and so we get

$$t - t_A = \frac{mc^2}{P_l} \frac{1}{2\tilde{r} + \tilde{a}} \beta \quad \text{(non-relativistic limit)}. \tag{11.20}$$

This is shown as the line labelled ‘classical’ in figure 11.7.

How far does the sail travel before it achieves a certain velocity? Using

$$\frac{d\beta}{dt} = \frac{d\beta}{dx}\frac{dx}{dt} = \frac{d\beta}{dx}\beta c \tag{11.21}$$

we can rewrite equation 11.16 in terms of $d\beta/dx$ and then integrate. With $x = x_A$ when $\beta = 0$ we get

$$x - x_A = \frac{mc^3}{3P_l} \frac{1}{2\tilde{r} + \tilde{a}} \left[1 + \frac{(2\beta - 1)}{\gamma(1 - \beta)^2} \right]. \tag{11.22}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

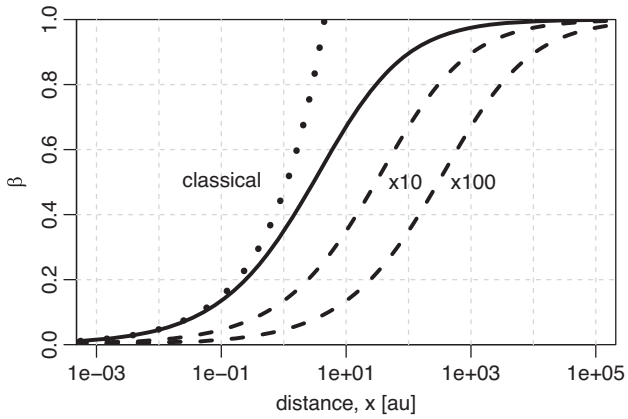


Fig. 11.8 As figure 11.7 but now showing the variation of the velocity with distance travelled x (equation 11.19).

This is plotted in figure 11.8, again for a 10 g sail spacecraft powered by a 100 GW laser. By the time the sail has reached a velocity of $0.1\,c$, the sail has travelled 7.8 million km (0.05 au), and when it has reached $0.2\,c$, it is 37 million km (0.24 au) away.

In the non-relativistic limit of $\beta \ll 1$, the term in square brackets in equation 11.22 becomes $3\beta^2/2$, and so we get

$$x - x_A = \frac{mc^3}{2P_l}\frac{1}{2\tilde{r}+\tilde{a}}\beta^2 \quad \text{(non-relativistic limit)}. \tag{11.23}$$

This is shown as the line labelled 'classical' figure 11.8.

We also need to know for how long we must run the laser in order to achieve a certain velocity. To do this we need to write equation 11.16 in terms of the retarded time T. As $T = t - x/c$, $dT = dt - dx/c$. With $\beta c = dx/dt$ this is

$$dT = dt(1-\beta). \tag{11.24}$$

As was already apparent from figure 11.4, time intervals between a pair of events at the sail, such as events A and B, are shorter in retarded time than in coordinate time because of the acceleration of the sail, and it therefore takes longer for the later photons to catch up with the sail. In terms of dT, equation 11.16 is

$$\frac{d\beta}{dT} = \frac{P_l}{mc^2}\frac{(1-\beta)}{\gamma}(2\tilde{r}+\tilde{a})\,. \tag{11.25}$$

This is shown as the dashed line in figure 11.6.

Integrating equation 11.25 with $\beta = 0$ at $T = T_A$, we find the velocity of the sail at retarded time T to be

$$\beta = \frac{(2q+1)^2-1}{(2q+1)^2+1}, \quad \text{where} \quad q = \frac{2\tilde{r}+\tilde{a}}{2mc^2}\int_{T_A}^{T} P_l(T^*)\,dT^*. \tag{11.26}$$

The integral is the total energy emitted by the laser between retarded times T_A and T, and therefore is the energy that will arrive at the sail between coordinate times t and t_A. This is the same as equation 11.7, but now with the energy replaced by a time integral of the power (over retarded time) and accounting for imperfect reflection. So we have generalized the case analyzed in section 11.1.1 of photons arriving simultaneously to photons arriving at

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

arbitrary times for constant or variable laser power. The integrated power is the total energy E_l provided by the laser; its variation with sail velocity for the perfect reflection case was shown in figure 11.2. If the power is constant, then $E_l = P_l(T - T_A)$, and so this figure also shows how the velocity of the sail varies with retarded time (that is, the velocity it will eventually achieve once all the photons have arrived).

As the sail recedes from the laser, its angular size will get smaller, so in reality an increasing amount of the transmitted laser power will be lost around the edge of the sail due to diffraction from a finite-sized transmitter. We will examine this in section 11.2.1. We could accommodate this through an appropriate variation of the power with distance.

11.1.5 The power received by the sail

Consider now the power P'_s received by the sail in its rest frame. How does this relate to the power P_e provided by the laser in its rest frame? In the laser rest frame, photons of frequency ν_e are emitted at constant rate $\dot{n}_l$, so $P_e = \dot{n}_l h\nu_e$. The power received by the sail in its rest frame can likewise be written $P'_s = \dot{n}'_s h\nu'_e$ (all primed quantities are measured in the sail rest frame). The photons hit the sail when it is moving at velocity β, so the sail sees them Doppler shifted to frequency $\nu'_e = \nu_e/\psi$, with ψ defined in equation 9.7. The primed and unprimed labels are swapped with respect to the notation in chapter 9, because now we consider the unprimed frame as the source of the photons. As the laser is moving away from the sail, we have $\alpha = \pi$ and $\beta > 0$ so $\psi = \gamma(1+\beta)$, which is greater than 1; the photons are redshifted. The rate at which photons are received is also a frequency and so we also have $\dot{n}'_s = \dot{n}_l/\psi$. Putting these two terms together we get $P'_s = P_e/\psi^2$. Using equation 7.6 we can write this as

$$P'_s = P_e\left(\frac{1-\beta}{1+\beta}\right). \tag{11.27}$$

This is the power received by the sail in its rest frame when it is moving at velocity β. The factor on the right decreases monotonically from one to zero as β increases from zero to one. Thus the power received by the sail is less than the power emitted by the beam, and the larger the velocity, the less power is received.

Where has this power gone? It is not in the reflected photons, because we are not yet considering the reflection process; we are just looking at the power that arrives in the rest frame of the sail. It is also not due to any practical losses, because we are dealing with an idealized situation. Part of the explanation is that equation 11.27 gives the instantaneous power, not the total energy. Not all of the energy emitted per unit time in the laser rest frame arrives in the same interval time in the sail rest frame. Because of the motion of the sail, some of the energy is still in photons that have not yet arrived at the sail. This can be seen by referring once again to figure 11.4. This is a purely classical effect. On top of this comes the relativistic effects of shifting between rest frames. An exercise at the end of the chapter examines this a bit more.

Even once all of the photons have arrived at the sail, not all of the energy emitted by the laser as measured in the *laser* rest frame is received at the sail as measured in the *sail* rest frame. This is because energy is not conserved between different rest frames, as noted earlier in the context of the Doppler effect (section 11.1.1).

When we previously considered the energy exchange from within the laser rest frame, we saw that at larger velocities a larger fraction of the laser energy is put into the kinetic

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

energy of the sail (figure 11.3). In the present section, we considered the energy exchange from within the sail rest frame, and found that the larger the velocities, the lower the rate of energy transfer to the sail. This apparent contradiction is resolved by the fact that previously we were comparing an integrated quantity (energy) in a common rest frame, whereas here we are comparing a differential one (power) between two different rest frames.

11.2 Optics and arrays

11.2.1 Diffraction

Solar sails are propelled by an isotropically emitting source, namely the Sun or other star. The wavefront from a star expands spherically, and when it reaches the spacecraft, the wavefront is close to flat with uniform intensity. But when we propel the sail with a directed source like a laser, we must consider diffraction. Even if we create a perfectly collimated beam, the fact that it passes through an aperture of finite size leads to an unavoidable spreading of the beam due to diffraction. Consequently, some of the laser power may pass around the sides of the sail. This is still the case when we use a lens or mirror to focus the beam to a point on the sail. What in geometric optics would have been a point will be extended. This can be understood in Kirchhoff's diffraction theory as the result of interference from a continuum of sources in the plane of the emitting aperture.

Consider a circular aperture – the laser source – of diameter d_t transmitting monochromatic light of wavelength λ (figure 11.9). The interference pattern on a surface (the sail) a distance r from the aperture can be described by Fraunhofer diffraction (far-field diffraction). For a circular aperture, the variation in intensity (power per unit area) on a plane perpendicular to the direction of propagation is

$$I(u) = I_0 \left(\frac{2J_1(u)}{u} \right)^2, \quad \text{where} \quad u = \pi \alpha \frac{d_t}{\lambda}, \tag{11.28}$$

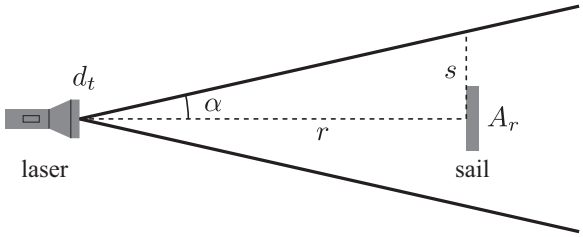
α is the angle from the direction of propagation, I_0 is the intensity at the centre, and $J_1(u)$ is a first-order Bessel function. This is shown in figure 11.10. For this function, 84% of the total power is contained between between the centre and the first minimum, a region known as the *Airy disk*. The first minimum is at $u = 1.22\pi$, so the angular radius of the Airy disk is $1.22\lambda/d_t$. A total of 91% of the power is contained within the second minimum at $u = 2.23\pi$, and 94% within the third minimum at $u = 3.24\pi$. The intensity at the centre of the Airy disk is

$$I_0 = \frac{P_t A_t}{\lambda^2 r^2} \tag{11.29}$$

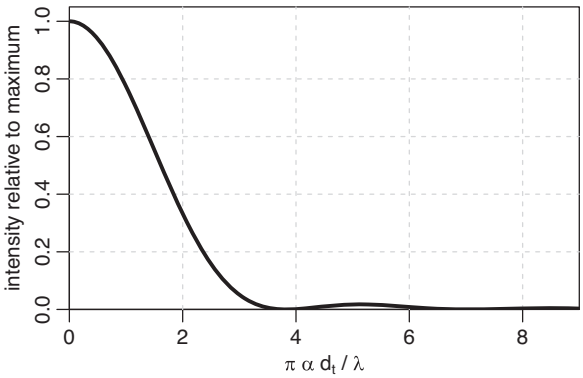
where P_t is the power emerging from the aperture, $A_t = \pi d_t^2/4$ is the area of the aperture, and r is the distance from the aperture. (We'll see in section 15.2 how the dependencies in equation 11.29 arise.)

The Airy disk is conventionally taken as the size of the beam at the receiver. For a 1 m telescope transmitting light at 1 μm for example, we get $\alpha = 1.22\,\mu\text{rad}$, or $0.25''$. For other shaped apertures, the angular size of the region containing a fixed fraction of the total power

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012



The diffraction of a light source through an aperture of diameter d_t. From diffraction theory, $\alpha \sim \lambda/d_t$, where λ is the wavelength of the light, so at the sail (of area A_r) the beam has a characteristic radius of $s \sim \lambda r/d_t$. **Fig. 11.9**



Diffraction pattern far from a circular aperture of diameter d_t with light of wavelength λ. This shows how the intensity varies as a function of angular distance α from the axis in the Fraunhofer diffraction regime. **Fig. 11.10**

still varies as λ/d_t but with a different numerical constant. Writing this constant as a, the *diffraction angle* can be written more generally as

$$\alpha = a\frac{\lambda}{d_t}, \tag{11.30}$$

with $a = 1.22$ for the first minimum of the Airy disk. The physical radius of the beam at a distance r from the source is $s = \alpha r$ from geometry when $\alpha \ll 1$. As an example, for a 10 m diameter telescope transmitting at $\lambda = 1\,\mu$m, then at a distance of $r = 10^6$ km the Airy disk on the sail has a radius of 122 m, and thus a diameter of 244 m. In the previous section, we saw that when accelerating a 10 g sail with a 100 GW laser, it will travel 7.8 million km by the time is reaches a velocity of 0.1 c. At this distance the Airy disk has a diameter of 1.9 km. This is much larger than the sail, which at this mass is unlikely to be much larger than a few metres. Thus almost all of the light would be lost around the edge of the sail by diffraction, meaning the sail would never have attained this velocity by this distance in the first place. To reduce these diffraction losses, we need to use a much larger telescope aperture. When the aperture is 1 km in size, for example, then $\alpha = 0.25$ mas and the Airy disk at 7.8 million km has a diameter of only 19 m, although this is probably still much larger than the sail.

The above discussion has assumed that the sail is stationary relative to the laser/telescope. When in motion, the diffraction angle is increased by aberration, as given by equation 9.5.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

When $\beta = 0.1$, the diffraction angle is 11% larger and the solid angle 22% larger, and this will spread out the laser power even more.

It is clear that we need very large transmitting optics to propel laser sails, and in the next section we'll look at how to achieve this. Even then, the transmitting aperture and sail will likely not be large enough for the sail to intercept all of the Airy disk at all distances during the acceleration. Limits on laser power mean the sail won't reach its final velocity before there is significant loss of power around the edge of the sail. This needs to be taken into account in a more detailed treatment of the sail kinematics (see exercises).

We discussed in the previous section that much of the laser power incident on the sail is reflected back to the laser. This could be reflected back again to the sail, thereby increasing the power or saving energy. The problem with this photon recycling is that the beam reflected by the sail is also diffracted due to the finite size of the sail. Because the sail is much smaller than the laser aperture, this diffraction is very large, so relatively little power returns to the laser, and even less makes it back to the sail again. This can be repeated, but every reflection dilutes the beam significantly.

11.2.2 Aperture synthesis

Building a kilometre-sized optical telescope would be extremely challenging, and almost impossible on the Earth's surface given the need to support and steer it. However, it is not necessary to build an individual telescope structure in order to achieve a large effective aperture with a narrow diffraction pattern. If we take a large telescope mirror and remove small parts of it, this changes the detailed shape of the diffraction pattern, but it does not significantly change the width of the central part of the pattern that contains most of the energy. A large telescope mirror with bits missing is equivalent to a number of smaller telescopes with the appropriate surface shapes and relative positions. This is the idea of aperture synthesis, shown in the left panel of figure 11.11. This allows a large effective aperture to be achieved without having to build a large single telescope. The 'baseline' is the separation between any two individual telescopes, and the largest baseline is equal to the diameter of the original monolithic telescope. The transmitting area (or collecting area, if it is being used as a receiver) of this effective aperture is reduced, but for many applications this is not a limitation.

Telescope arrays are commonly used in astronomy, so let us first consider using the array as a receiver to observe a distant object (incoming plane wavefront). The array shown in figure 11.11 (left panel) has the individual telescopes in the physical locations they would have in the filled aperture. This is impractical for a large array on the Earth's surface, because to point it in different directions we would have to lift individual telescopes far above the ground. However, the purpose of the precise shape of a mirror is to ensure that all of the wavefront reflected off the different parts of the mirror travel the same distance and so maintain a common phase at the focus. We can instead place the individual telescopes in a common plane, and then (somehow) remove the phase shifts that arise because different portions of the wavefront that went to different telescopes had to travel different distances (figure 11.11, right panel). These phase shifts are called the 'geometric delay'. By varying the phase shifts, the effective aperture is sensitive to different directions, equivalent to pointing it in different directions. This is the idea of an astronomical interferometer. The resulting diffraction pattern (also for a filled aperture) is often referred to as the *point spread function* (PSF).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

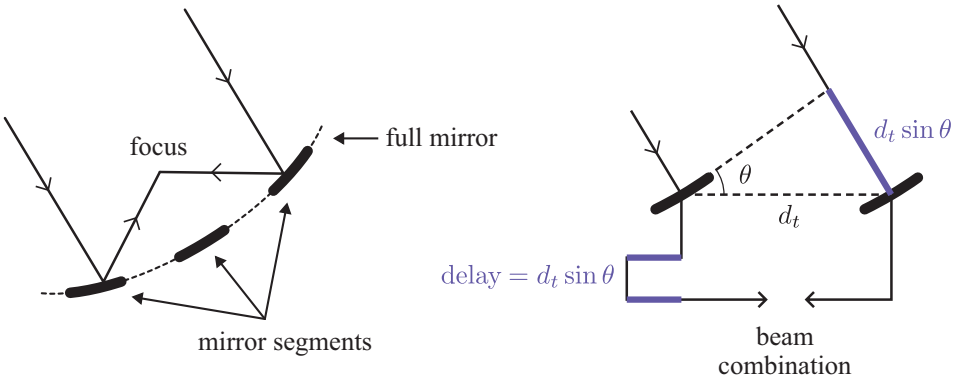


Fig. 11.11

Aperture synthesis. Left: Parts of the the primary mirror of a telescope can be removed to leave just some mirror segments. This reduces the amount of light gathered and changes the diffraction pattern, but the telescope still has an effective aperture diameter equal to the original complete telescope. Right: This same-sized aperture can be synthesized with the mirror segments in a common plane (e.g. the surface of the Earth) by introducing appropriate optical path delays between the wavefronts received at the different segments to ensure that the wavefronts are in phase when combined.

How can we introduce the phase shifts? Astronomical interferometry was first developed most extensively at radio wavelengths, because at these longer wavelengths we can detect both the amplitude and the phase of electromagnetic waves and record them electronically. This allows the signals from the different individual telescopes to be combined electronically too. As the signals are stored on a computer, the combination does not even have to be done in real time. This has enabled the combination of signals from radio telescopes separated by thousands of kilometres and thus the synthesis of apertures on this scale, in a process known as very long baseline interferometry (VLBI). It was this technology that allowed the first images of black holes to be obtained with the Event Horizon Telescope in 2019, for example.

At optical wavelengths, the direct electronic detection of phase is currently impossible, so astronomers instead remove the phase shifts by making the light that arrives at the first telescopes travel a longer path, before optically combining the light from all the telescopes, as shown in the right panel of figure 11.11. By varying the length of these delay lines – at the same precision required for a telescope mirror to keep wavefronts in phase – the effective aperture can be pointed in different directions. (The individual telescopes would be pointed too.)

In principle, this same process can be used backwards to transmit power from a powerful light source through an optical telescope array. This would produce an outgoing wavefront with a common phase that acts as though it were a beam focused from a single large telescope. In practice this would be very difficult, on account of the extremely high light intensity that has to be transmitted through the delay line optics. It is also not possible to produce an individual laser with anything close to the required power. We would instead need to feed the telescopes with separate lasers, and somehow ensure that the lasers have the appropriate phase offsets needed to synthesize the larger aperture. We look at how to do this in section 11.3 below.

With a phase-corrected array of telescopes, the size of the central part of the PSF it produces is of order λ/d_t, where d_t is the largest baseline in the array. However, the amount of power in that region is less than in the case of an Airy disk produced by a filled aperture. This

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

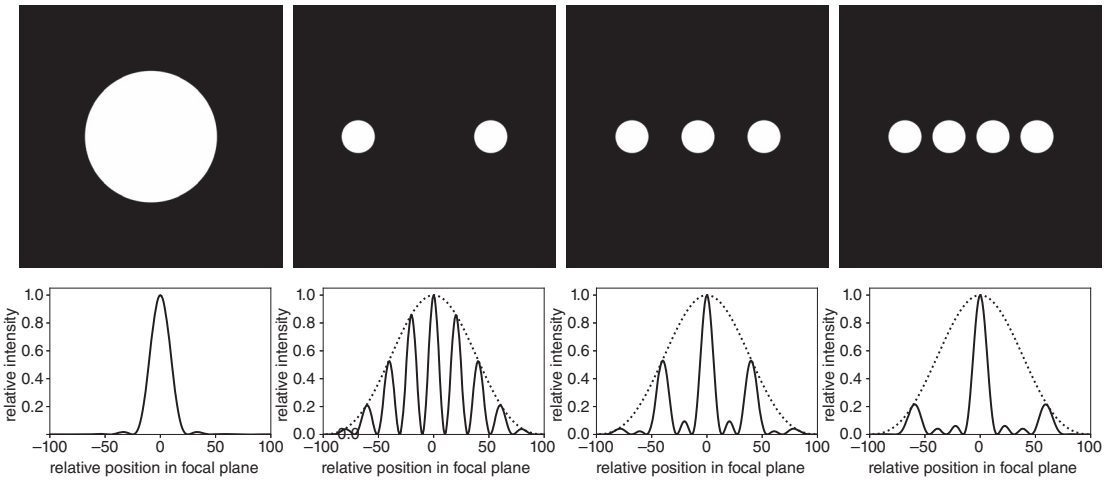


Fig. 11.12 Four different apertures (top row) and their corresponding PSF (bottom row). The white circles in the top row show the size and location of individual telescope apertures that are used to synthesize an aperture. The largest baseline is the same in all cases (equal to the diameter of the single large aperture in the left panel). The bottom rows shows the variation in intensity of a horizontal cut through the middle of the PSFs. The dashed line in these panels shows the PSF of a single small aperture for comparison. The scale on the horizontal axis is common to all panels, but in arbitrary units, and their vertical axes have been scaled to the maximum intensity in each case.

can be seen in figure 11.12, which shows four different apertures that have the same largest baseline d_t, but a different number of telescopes and thus a different filling fraction of the synthesized aperture. (In practice the array would extend in the other direction too; a linear array is shown for simplicity.) The first PSF on the left is for the filled aperture of diameter d_t. This is the Airy pattern mentioned earlier. The second PSF comprises two smaller apertures. The central peak of this PSF is narrower than the Airy disk for the filled aperture, but a larger fraction of the power is now in the so-called side lobes. As we increase the number of same-sized apertures to three and then four at constant baseline, we see that a larger proportion of the power (total area under the curve) is concentrated into the central peak of the PSF. The dashed line is the PSF for a single small aperture, which shows that the envelope of the synthesized PSF is determined by the individual aperture size.[4]

For most situations it turns out that the array can deliver more power to a sail when we use a compact array, with individual telescopes as close to each other as possible, rather than a sparse array with a larger effective aperture (so narrower central peak) but more power put into the side lobes.

11.2.3 Atmospheric distortion and adaptive optics

A significant complicating factor for an optical array on the Earth's surface is the atmosphere through which the light must pass. Spatial and temporal variations in the refractive index of the air – caused by temperature variations, wind, and turbulence – distort the wavefront. This introduces phase differences across the wavefront on a range of length and timescales.

[4] More precisely, the PSF of the synthesized aperture is the product of the Fourier transform of an individual telescope aperture with the Fourier transform of the array (represented as a set of delta functions).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

In astronomy, this wavefront distortion is referred to as 'seeing'. Without correction, seeing limits the size of the PSF in the optical and near-infrared to be no less than about $1''$ (to within factors of a few). This is a thousand times larger than the diffraction limit we want to achieve.

Seeing not only limits the angular resolution, but also spreads faint light from a point source across a larger area, mixing it with diffuse background light, thereby reducing the sensitivity of the observations. The seeing varies with wavelength (better at shorter optical wavelengths) and altitude (better on mountains). But with a seeing of $1''$, then at 1 μm, a telescope with an aperture larger than 0.25 m is seeing-limited rather than diffraction-limited. This applies equally well to an aperture synthesis telescope. If we cannot overcome the seeing, the spot size of laser power sent through the array will still be of the order of $1''$, and not reduced by the large extent of the array.

The problem of seeing is addressed in astronomy by measuring the distortion of the wavefront and using this information to adjust the shape of the mirrors to undo the distortion. This is known as adaptive optics, the principle of which is summarized in figure 11.13. The wavefront can be modelled as a contiguous set of small patches that have become tilted with respect to one another by the atmosphere. By taking some of the light from the observed source, the tilt of each patch can be determined. This can be done via intensity measurements. Instead of the target source itself – which may be too extended or too faint to provide enough light for wavefront monitoring – we can instead use a bright star that lies a small angular distance from the target, a so-called natural guide star. In either case, these tilt measurements provide a map of the distortions across the pupil (aperture) of the telescope. A deformable mirror in the pupil plane that has subapertures corresponding to the measured patches is then adjusted to undo the tilts. The result is that the wavefront reflected off this mirror is restored to being flat.

For this correction procedure to work, each patch must be small enough such that the wavefront distortion can be approximated by a simple tilt about two axes. This patch size is known as the Fried parameter r_0, and is of the order of a few tens of centimetres, depending on wavelength and atmospheric conditions. The wavefront must also be measured, and the mirror adjusted, on a timescale that is shorter than the variation timescale of the wavefront. As the turbulence in the atmosphere evolves relatively slowly, this timescale is set primarily by the time it takes air to cross the patch, and so is of order r_0/v, where v is the wind speed. Timescales are typically a few milliseconds, so mirror adjustments need to be made several hundred times per second.

Adaptive optics for near-infrared wavelengths is routinely employed on large ground-based telescopes, such the four 8 m telescopes of the ESO's VLT in Chile. Building such a large primary mirror that can be adjusted at high frequencies would be challenging. But as large telescopes invariably employ smaller mirrors to help focus the light, one of these is usually the deformable one.[5] The VLT secondary mirrors have a diameter of 1.12 m, and one of them has 1170 actuators that adapt its shape at high frequency.

When individual telescopes are connected to make an interferometric array (as can be done at the VLT Interferometer), adaptive optics can be applied to each telescope independently. With the wavefront from each telescope beam corrected in this way, and the appropriate

[5] The primary mirror is often adjustable too, but at a much lower frequency, foremost to compensate for the distorting effects of gravity when this large thin mirror is tilted.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

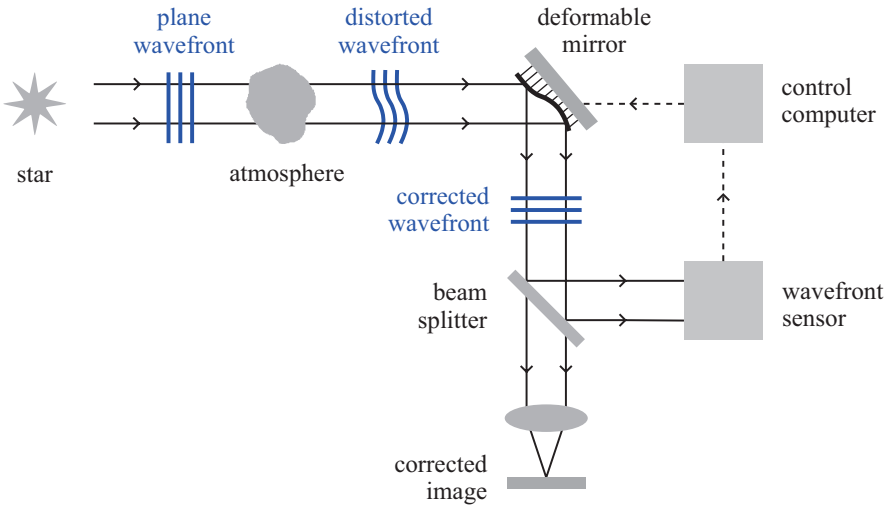


Fig. 11.13 A schematic outline of an astronomical adaptive optics system. The wavefront from a star, distorted by the atmosphere, enters the telescope and reflects off a deformable mirror. The wavefront is then divided, with a portion of the light going to a wavefront sensor that measures the distortions, and the other portion going to the image plane. The information from the wavefront sensor is used by a control computer to adjust the form of the deformable mirror to remove the wavefront distortions, so that the corrected wavefront can be imaged in a camera. The sensing and mirror correction typically takes places a few hundred times per second.

optical delays introduced for each beam to synthesize an aperture (figure 11.11), the array can achieve a PSF size that is determined by the diffraction limit of the array, rather than by the seeing or the individual mirrors.

The same issues apply to transmitting light from a laser array to a sail. An outgoing laser beam will have its wavefront distorted by the atmosphere just as incoming starlight does. If we want the laser array to be diffraction-limited to achieve the narrowest possible beam on the sail, then we need to correct these distortions. This could again be done with deformable mirrors. To measure the distortions we again need a coherent light source above the atmosphere that is imaged by the adaptive optics system. Natural guide stars could probably not be used because they would be much fainter than any backscattered laser light (and also fainter than the scattered sunlight if the laser is used during the day).

As the laser light leaving the telescope is designed to be coherent, one solution might be to use laser light reflected off the sail (recycled photons) for the adaptive optics corrections. The laser is powerful, and although the reflected photon beam is strongly diffracted by the small size of the sail, there would be enough photons for the corrections. Unfortunately, though, by the time the photons have travelled up from the telescope to the sail and back again, the atmosphere will have changed. If the atmospheric timescale is 1 ms, then the maximum distance to the sail permitted for this approach would work is $c \times 10^{-3}/2 = 150\,\text{km}$. This is barely above the atmosphere, yet we saw in section 11.1.4 that the laser needs to propel the sail over millions of kilometres. So we can't use this reflected light for the corrections either.

Some of the outgoing laser light will be reflected back to the telescope by the atmosphere, and this could potentially be used for the wavefront correction. This is rather complicated, however, as the reflections occur across a range of altitudes, so at best only a partial correction could be achieved.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

A better solution to this problem is to place a satellite above the atmosphere, close to the path of the laser, that shines a laser back at the telescope array, thus acting as a bright artificial guide star. This has already been explored as a possibility for use by astronomical telescopes. Such a beacon satellite needs to be in a relatively high orbit so that its downward beam – which is conically expanding due to diffraction – passes through the same portion of the atmosphere as the upward-propagating beam from the telescope array (which is conically expanding in the other direction). Without this, the atmospheric distortion cannot be properly corrected, a problem with laser guide stars in astronomy known as focal anisoplanatism, or the cone effect. A satellite altitude of at least 160 000 km has been suggested, which is over four times the geostationary altitude. The beacon's orbit has to be set according to the location of the telescope array and the direction to the target (star or sail). As the beacon needs to be very close to the line of sight between telescope and target, yet is moving, it should be placed in a highly eccentric orbit and used near apogee, where it is moving the slowest. Even then it would be usable only for short time periods. A beacon in a non-Keplerian orbit, achievable with an ion engine, or the deployment of multiple beacons, could be considered.

11.3 Lasers

11.3.1 How much power?

To accelerate a 10 g sail to $0.1\,c$ requires that it receive 4.5×10^{12} J or 1.3 GWh of energy. As we have seen, the laser has to provide more energy than this because most of the energy remains in the reflected photons and is not transferred to the sail. From figure 11.3 we see that the overall efficiency to reach $0.1\,c$ is only 10%, so in this case the laser must provide 13 GWh of energy. If this energy is transmitted at 1 μm by a telescope array with an effective diameter of 1 km, then a laser spot focused on the sail when it is 0.01 au from the array has spread by diffraction to an Airy disk diameter of 3.7 m (a bit more when we include aberration). This is the size of a 10 g sail with a lightness number of 1.6 (equation 10.10), which is large but plausible with newly identified materials (discussed below). Thus in order to avoid a significant loss of laser energy due to its beam spilling out around the sides of the sail, we must transfer all of the energy to the sail when it lies within 0.01 au. Assuming a constant acceleration up to $0.1\,c$ over this distance, this means we must provide the energy within about $0.01\text{au}/(0.05\,c) = 100$ s (actually about 5% less due to the light travel time, but we are just doing an order of magnitude calculation here). This implies we need a laser power of 13 GWh/100 s, or 500 GW.

Compared to currently available lasers (discussed below), 500 GW is an enormous power for a laser that must operate for this duration. Furthermore, it corresponds to an intensity on the sail that surpasses what current materials can withstand (see section 11.4.1 below). With a laser of only half the power we could still achieve the same final velocity for the same sail if we illuminate it for twice as long, which will also lower the intensity. However, to keep the same diffraction losses we would have to double the size of the telescope aperture, in this case to 2 km.

In practice, both laser power and array size will be limited, so we will have little choice but to accept a longer acceleration time, and therefore diffraction losses, and therefore an even

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

longer acceleration time. The energy efficiency will go down, and storing enough energy – and then delivering it fast enough – is also a major challenge.

The sail mass and size are critical to the scaling of the problem. If the sail is twice as massive, it requires twice as much energy to get to the same final speed. Assuming fixed mass per unit area, then with twice the mass, the diameter is $\sqrt{2}$ times larger. So using the same telescope array (same PSF size), the sail can now travel $\sqrt{2}$ times as far before suffering the same diffraction losses, and thus has a travel time that is $\sqrt{2}$ longer. But as power is energy divided by time, this requires $2/\sqrt{2} = \sqrt{2}$ times as much power. In other words, for a given telescope array, a more massive sail also requires more laser power, not just more energy. This favours small sails. However, the sail must also carry a payload that adds mass but does not contribute to the propulsion, so a minimum-sized sail is required for a given payload mass. In practice there is a complicated trade-off between array size, sail size, laser power, and diffraction losses, as well as total energy required.

11.3.2 Arrays of lasers

Lasers of tens or hundreds of gigawatts of power that can run for extended periods of time do not yet exist. Some of the most powerful lasers available are those used in inertial confinement fusion, discussed in section 8.5. Each of the lasers in the National Ignition Facility, for example, has a power of about 2.5 TW, but delivers this for only a few nanoseconds. Perhaps the most powerful long-duration laser was the military MIRACL laser, able to deliver 1 MW for about a minute. Others have powers closer to tens of kilowatts. The input power for such lasers is many times larger than their laser output, and power is additionally required for water cooling. In any case, a large number of individual lasers will have to be combined to achieve tens or hundreds of gigawatts.

There are many different laser technologies operating at a wide range of wavelengths. An Nd:YAG laser, for example, is typically used to lase at 1064 nm. There are other transitions that can be used, and frequency doubling would lower the wavelength, which is beneficial for reducing the diffraction. Solid state diode lasers also offer many possibilities.

As discussed earlier, we need to synthesize a large transmitting aperture using multiple smaller telescopes (section 11.2.2). To achieve this, we could split the light from each laser across all of the telescopes, such that the entire array provides a coherent beam for each laser separately, and these beams then overlap incoherently on the sail. But this is impractical, because it would require a lot of optics in order to split and geometrically delay the light from each of the lasers (tens of thousands) into each of the telescopes (thousands).

A more practical approach is for each laser to feed only one telescope. In this case we need to ensure that all the lasers are in phase, and we have to introduce the correct geometric delay for all the different lasers that feed a given telescope. All the beams from the different individual telescopes are then coherent with one another and the outgoing wavefront has the diffraction limit of the array. This approach poses significant challenges too. Getting and keeping all the different lasers in phase is difficult, because the frequency and phase of a laser can drift over time due to numerous effects, including temperature variations, acoustic vibrations, and noise in the electric current driving the laser. The different lasers must therefore be locked to the same frequency and phase. This can be done using a common external

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

reference, together with one of several active methods that rely on feedback from the laser to monitor changes and make corrections.

In all of this we have been assuming that the laser is ultraviolet, visual, or near-infrared. Longer wavelength sources, such as masers, bring simplifications in the aperture synthesis. Specifically, coherence could then be achieved electronically instead of using optical delay lines, in a so-called phased array antenna, which we will examine more in section 15.8.3. However, to keep diffraction losses within reasonable limits, the apertures would need to be two or three orders of magnitude larger, which seems to rule them out. Phased arrays in the optical are possible by integrating the delay lines in solid state without any moving parts. If these could be scaled up over very large areas to deliver the huge amounts of laser power required, this could do away with the need for large individual telescopes.

11.4 Sail structure

We have looked at the energy and power requirements of a sail and at how we might get this power to the sail by means of a laser and telescope array. We turn now to the properties of the sail itself.

11.4.1 Materials

In the context of solar sails (section 10.5), we identified that the ideal sail material has low mass per unit surface area, high reflectance (equation 10.21), and low absorptance and/or high melting temperature (equation 10.25b). These requirements cannot be optimized independently of one another. We saw that a very thin metallic sail will eventually become transparent, for example. A sail propelled by a laser will be exposed to much larger intensities than a solar sail. In section 11.3.1 we considered a sail of 3.7 m diameter, and a laser power of around 500 GW. Even if we consider a more modest 100 GW, this still corresponds to an intensity on the sail of $9\,\mathrm{GW\,m^{-2}}$. In comparison, at a distance of 10 solar radii, a solar sail is exposed to an intensity of less than $1\,\mathrm{MW\,m^{-2}}$. Laser sailing must pay much closer attention to the thermal properties of the spacecraft.

The Stefan–Boltzmann law in equation 10.25b tells us that the equilibrium temperature of a sail increases only slowly with the incident intensity, namely to the one-quarter power. Nonetheless, with a laser intensity of $9\,\mathrm{GW\,m^{-2}}$, and a sail absorptance of 0.01 and emissivity of 0.1, this equation predicts an equilibrium temperature of 9400 K. This is well above the highest melting temperature of any known material, which is around 4000 K (diamond 3800 K; tantalum hafnium carbide alloy 4200 K). Even with unit emissivity, the absorptance would have to be below 0.003 to get down to these temperatures (with any factor decrease in the emissivity having to be offset by the same factor decrease in the absorptance). In practice, the maximum operating temperature would have to be lower, because absorptance tends to increase with temperature, and materials may weaken well before they melt. In any case, the absorptance will never be zero, so the sail should have a high reflectance at the wavelength of the incident laser beam – optical or near-infrared – yet a high emissivity at longer wavelengths corresponding to its equilibrium temperature. By choosing the right material(s), the reflection and emission can be optimized for these different wavelengths.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

Although metals are highly reflective, they have neither very low absorptances nor high melting temperatures. Metallized films like those considered for solar sails are therefore inappropriate. For laser sail materials we need to turn to semiconductors, dielectrics, and composite materials. Zirconium dioxide (ZrO_2), for example, has a high emissivity of 0.95 and melting temperature of 2990 K. Its refractive index n at 1 μm is around 2.13, and so a thick piece of material has a reflectance of $[(n-1)/(n+1)]^2 = 0.13$. Germanium has a higher refractive index of around 4.1, and therefore a higher reflectance of 0.37.

The reflectance for monochromatic light of wavelength λ can be increased by using a thin film of dielectric material with thickness $\lambda/4n$, a result of the constructive interference of the beams reflected from the front and the inside back of the film. For such a 'quarter-wave plate' the reflectance increases to $[(n^2-1)/(n^2+1)]^2$, which for ZrO_2 is 0.41. Assuming zero transmittance, this implies an absorptance of 0.59. Using equation 10.25b, we conclude that a quarter-wave plate of ZrO_2 could tolerate an intensity of about 7 MW m^{-2} (with $\epsilon_b = 0.0$, so emission is only from the front). This is still far below the intensities considered in the previous section.

By using many layers of material, we can increase the reflectivity further, albeit at the cost of increased mass. Hein et al. (2017) cite several examples of multiple layers of dielectrics on metallized plastic films. A set of 15 layers of silicon dioxide (SiO_2) and titanium dioxide (TiO_2) dielectrics on a copper film achieves a reflectance of 0.9999294 (for light of wavelength 1.06 μm) and emissivity of 0.07. This can tolerate an intensity of nearly 0.2 GW m^{-2} at its maximum temperature of 1360 K.

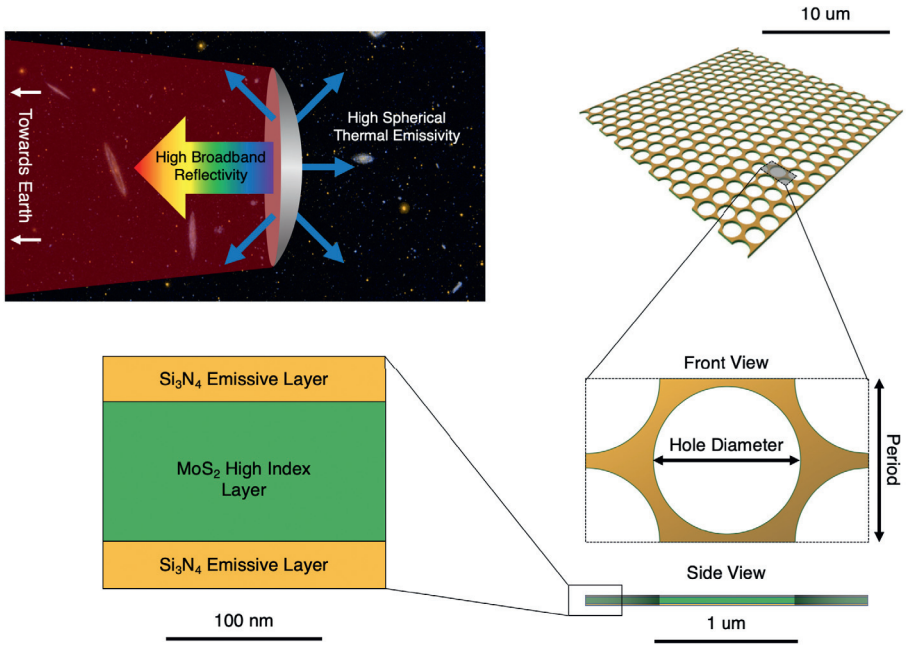
The problem with using quarter-wave plates in this way, however, is that they are tuned to a single wavelength, whereas the sail will see an increasingly redshifted laser beam as it accelerates. One solution to this might be to vary the wavelength of the laser during the acceleration (difficult in practice), or to accept a lower average reflectivity due to imperfect tuning of the sail to the range of incident wavelengths.

Two other materials with high reflectance and low absorptance that could be suitable sail materials are molybdenum disulphide (MoS_2) and silicon nitride (Si_3N_4). The former has a high refractive index of around 3.7, and single layer samples of this have been measured to have negligible absorptance at visible wavelengths. Their densities are not that low, however (5060 kg m^{-3} for MoS_2).

One way to reduce the mass of the material is to fabricate it as a two-dimensional photonic crystal. This is a layer of material with a regular grid of micron-sized holes that have a diameter not much less than their separation, so most of the structure is holes (figure 11.14). Not only does this reduce the mass, but because the length scale of the repeating structure is similar to the wavelength of the light, this actually increases the reflectance. This is due to the interference produced by a large number of holes, similar in concept to a diffraction grating. So far it has been possible to manufacture photonic crystals of MoS_2 many centimetres in size just 80 nm thick, with holes 1 μm in diameter separated by 1.14 μm, which achieves a reflectance of 0.90. The thermal emissivity is increased by adding a 5 nm layer of Si_3N_4 to each side.

Counter-intuitively, some materials with very low reflectance may be of interest if they remain strong enough even when extremely thin. One such material is graphene, which comprises a single layer of carbon atoms arranged in a hexagonal structure. It is therefore extremely light, with an area density of just 7.4×10^{-7} kg m^{-2}. Graphene normally has a very low reflectance and an absorptance of just a few per cent or less (and thus a high

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012



A two-dimensional photonic crystal. Credit: Reprinted with permission from Brewer et al. (2022), copyright 2022 American Chemical Society. Fig. 11.14

transmittance), although doping it with impurities might raise the absorptance to 0.40. With emissivities of $\epsilon_b = \epsilon_f = 0.95$ and adopting a maximum temperature of 3000 K, this would allow the sail to operate at intensities of up to 20 MW m^{-2}. Even with zero reflectance, a sail from this material would have a lightness number (equation 10.20) of 410, but both the doping and the payload would lower this value considerably.

Another form of carbon that is promising for laser sails is aerographite. This is a mesh of carbon tubes with an incredibly low density, just 0.2 kg m^{-3}, which is 14 000 times lighter than aluminium and six times lighter than air. If it can be made into a sail, it would provide an extremely high (pre-payload) lightness number. Its absorptance is near unity and it has a melting temperature around 3900 K. With unit emissivity, the maximum intensity it can tolerate is about 25 MW m^{-2}.

The maximum tolerable intensities for all these materials is much less than the GW m^{-2} scale intensities on a sail we computed at the beginning of this section, so major technological developments are still required. However, with very light materials, the mass and thus the lightness number of the spacecraft will in practice be determined by the mass of the payload and the size of the sail; the mass of the sail itself will hardly play a role. We can therefore make the sail much larger. This will permit a larger diffraction angle and thus a smaller telescope array, for a given power and illumination duration, which will lower the intensity on the sail. Alternatively, a longer illumination duration could be used for a given array size, thereby allowing the laser power to be decreased.

11.4.2 Shape and stability

So far we have implicitly been assuming that the sail is flat and circular, corresponding to the circular symmetry of the laser beam's diffraction pattern. A flat, specularly reflecting sail is

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

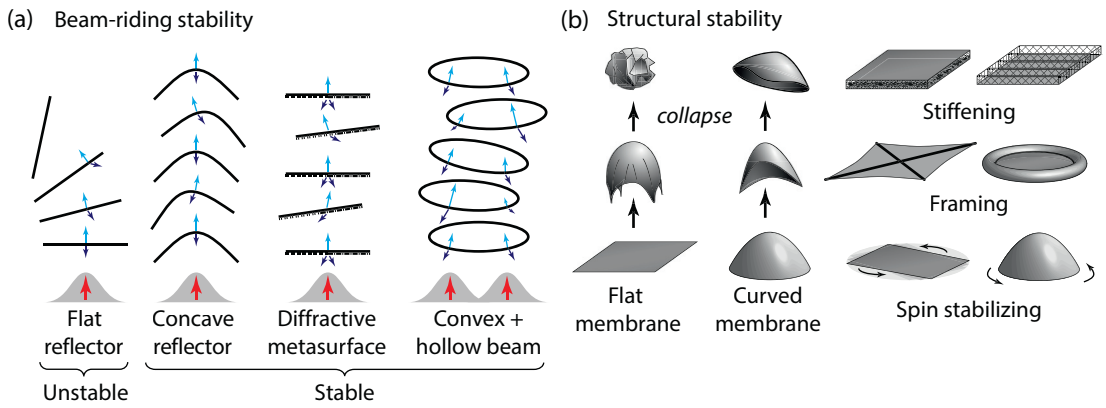


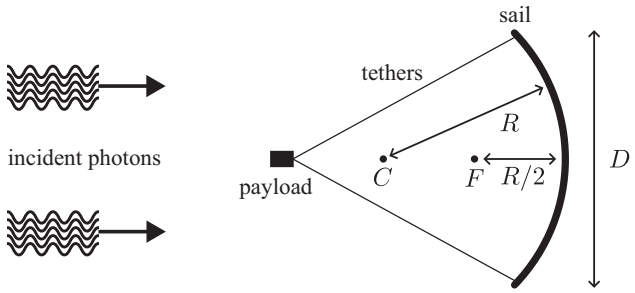
Fig. 11.15 Stability of lasers sails. Left: The orientation stability of different shaped sails when tilted slightly with respect to the beam direction. Right: The structural stability of various sails to the pressure of the laser beam. Credit: Gao et al. (2024), licensed under CC-BY-4.0, obtained from Nature Communications.

unstable, however. If the laser beam's intensity profile is not perfectly symmetric, or if the beam is not perfectly centred on the sail all the time, the resulting uneven photon pressure creates a torque on the sail and tilts it. This tilt not only changes the direction of travel of the sail, but would likely increase the tilt further and flip the sail out of control (figure 11.15). Some concave reflectors, in contrast, are stable to small variations in the beam intensity or centring. This includes conical, parabolic, and hyperbolic shaped structures. Diffusely reflecting surfaces can also be designed to give stability, even for flat surfaces. Convex structures, like a gas-inflated sail, can be stable if a 'hollow' light beam is used, that is, one in which the PSF is a circular ring with a minimum in the centre, rather than a Gaussian-like bump with a maximum in the centre. When the ring is similar in size to the sail, with the sail in the central intensity minimum, it provides stability against perturbations that try to shift or tilt the sail laterally.

Stability could also be achieved through feedback, but this would require additional sensing and control mechanisms acting at high frequency. Any feedback loop between the laser and the sail that adjusts the profile of the former in response to the orientation of the latter is infeasible given the large light travel time.

The above refers to the 'beam-riding' stability of the sail, namely its ability to remain on the beam and moving parallel to it. Another type of stability is structural, so that the sail retains its shape and does not tear. A flat sail would likely crumple when exposed to a Gaussian-like beam because the intensity would be much larger in the centre than at the edges (figure 11.15). To be stable against collapse, a flat sail would have to be able to resist very large stresses. A curved sail, in contrast, experiences much lower stresses. This is similar to the way in which parachutes or boat sails billow in the wind to minimize wind-induced stress. If I is the laser intensity on the sail, then for a spherical sail, the maximum tensile stress in the sail scales as IR/b, where R is the radius of curvature of the sail and b is its thickness (figure 11.16). Thus to minimize the stress for a sail material of a given thickness, we should minimize R, by having a maximally curved sail. However, to avoid the light incident on the sail's perimeter from reflecting back onto another part of the sail, we need $R > D/\sqrt{2}$, where D is the projected diameter of the curved sail. (Recall that the focus of a spherical reflector is at a distance $R/2$ from the surface of the reflector on its axis.) If we set $R = D$, then the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012



A spherical laser sail of diameter D and curvature radius R. C is the centre of curvature and F is the focus. To be stable the payload must be located at least a distance $R/2$ from the sail surface along the axis. The payload is shown here at $3R/2$. Fig. 11.16

surface area, and thus mass, is only 7% larger than a flat sail. Such concave sails gain stability at the cost of only a small mass penalty.

The sail must also carry a payload. From the point of view of mechanical stability, for a concave sail it is best to put the payload on the front (laser) side of the sail, attached to the sail's perimeter by tethers. For a spherical sail this is stable once the payload is beyond the focal point, that is, further than $R/2$ from the surface of the sail along the axis (figure 11.16). However, the payload is then fully exposed to the high-intensity laser beam. Another possibility might be to distribute the payload across the backside of the sail, but this would affect the stability.

A flat laser sail can be stable if it is made from a material that has anisotropic reflective properties. The idea here is to use nanophotonic elements to control the phase of light that is reflected from different points on the surface. It is possible to design this in such a way that if the sail tilts away from its nominal orientation, the reflections produce a restoring torque that acts to return the sail back to its nominal orientation. Sails that reflect light using diffraction gratings can likewise be designed to be self-stabilizing.

11.5 Lasers in space

From a construction perspective, it makes sense to place a large, powerful laser array on the surface of the Earth. This has several drawbacks, however. We have already looked at the problem of distortion of the wavefront by the atmosphere, which makes it difficult to produce a coherent beam from the array of apertures (section 11.2.3). The atmosphere will also absorb some of the laser energy, and this absorption itself can exacerbate the atmospheric turbulence. Absorption increases with the air mass through which the beam has to travel, which in a plane-parallel approximation varies as $\sec\theta$, where θ is the zenith angle. So at $\theta = 60°$ the absorption is twice that at the zenith. Aircraft, birds, and orbiting satellites may also be damaged or destroyed by passing through the high-intensity beam.

Another disadvantage of an Earth-bound laser is that it can access only half of the sky at any one time, and less than a quarter at low zenith angles. Even over the course of a year, a laser placed anywhere on the Earth other than on the equator never sees all of the sky. As the laser, sail, and target star must be more or less colinear, this limits the sets of stars that can be targeted with a single laser. The Earth's rotation also places a limit on how long the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

laser can be used to propel a given sail. If the required illumination time is more than a few hours, then either there have to be daily breaks in the transmission, or additional lasers have to be placed at different longitudes around the Earth. Target stars nearer to the pole, such as Alpha Cen (declination −63°), offer the possibility of continuous viewing from a high latitude.

Placing the laser array in space gives more flexible access to the sky. This also overcomes the problem of atmospheric turbulence, obviating the need for adaptive optics (although we still need to phase lock the different lasers to achieve a coherent synthesized beam). We would want to place the array in an orbit far from the Earth's surface, because in a low Earth orbit the Earth itself still blocks half the sky, and there would still be many satellites in higher orbits that might cross the beam.

The big drawback of lasers in space is how to power them. If the laser power is as low as a few gigawatts, nuclear fission reactors may be feasible. Alternatively, the power could be beamed from the Earth's surface at either visible or microwave frequencies. This beam would of course suffer from diffraction, but the distance would only be thousands of kilometres, and so the diffraction minimal. Power could only be transmitted for the time that the array is above the horizon as seen from the power source, a duration that depends on the altitude and thus the orbital period of the array.

An alternative approach is to use the power of the Sun. With an intensity of $1.36\,\mathrm{kW\,m^{-2}}$ at the Earth's distance, we require a collecting area of $0.74\,\mathrm{km^2}$ for each gigawatt of input power required, neglecting conversion losses. When covered with photovoltaics, the sunlight could be converted to electricity and this used to power a laser.

A major non-technical problem with a space-based laser is that it could be repurposed as a weapon to hit either satellites or targets on the Earth's surface. This dual use could make it very hard to get political approval for laser-based propulsion. An Earth-based laser is almost as problematic, because it could be used to destroy satellites, or even target locations on the Earth's surface when used in conjunction with reflectors in space. One solution to this objection is to place the laser on the dark side of the Moon, where it never sees the Earth. The Moon's lack of atmosphere and slow rotation overcome many of the problems with an Earth-based laser, and we could imagine having very large photovoltaic collectors to power it.

A laser can be avoided entirely by using a large mirror in space to collect sunlight and focus it directly onto the sail to propel it. The big challenge of this approach is that in order to transmit a focused collimated beam towards the sail – and not just reflect light diffusely – the mirror has to have a precise surface over its entire area, which is $74\,\mathrm{km^2}$ for 100 GW. This may be very hard to achieve.

The nearer we bring such a big mirror to the Sun, the smaller it can be for a given power. It would need a total area of 'only' 0.74 km when operating 0.1 au from the Sun (for 100 GW), and could even be segmented. The mirror could be made of very thin, reflective material, and in fact could be quite similar to a solar sail. Indeed, if we make its lightness number large enough, we could even use it as a sail, and by tilting, get it to spiral closer to the Sun without having to use rockets (see section 10.3.2). Naturally, when we then use it as a reflector to propel our sailcraft, the mirror itself will be pushed away from the Sun, although as it is significantly more massive that the sailcraft, its motion would be much less. Once we have finished propelling the sailcraft, we can again use the mirror as a tilted solar sail to return it to its initial position close to the Sun, ready to propel another sailcraft.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

11.6 Laser sail missions

Although sails propelled by solar photons have been successfully tested in space, laser sails have not yet been deployed. A number of concepts have been developed, however, both for interplanetary and interstellar missions.

There are many trade-offs and optimizations to make when designing a laser sail mission. For example, to accelerate a given sail to a specified velocity, we need a certain amount of energy. This could be provided either by a high-power laser over a short period of time, or by a lower-power laser over a longer period of time. A short period of time means the sail is nearer to the laser source on average during the acceleration, and so diffraction losses are kept small. But the larger power means a larger intensity on the sail, and so a higher temperature that is more demanding of the sail materials. Higher-power lasers are also more difficult and expensive to build, and to retain cool during operation. The longer period of acceleration required for a low-power laser, on the other hand, means either there are larger diffraction losses (and so more energy in total is required), or we need a larger laser array to reduce diffraction, which also comes with an expense.

With these kinds of trade-off in mind, let us look at a few specific mission designs.

11.6.1 High-mass payloads

Forward (1984) looked into using very large sails to carry payloads of many tonnes. One of these concepts is explored in detail by Frisbee (2004), who proposes a sail with a diameter of 32 km and a mass of 78 t carrying a payload of 100 t. This corresponds to an average areal mass density of the spacecraft of $2.3 \times 10^{-4}\,\mathrm{kg\,m^{-2}}$. It uses an aluminium layer just 16 nm thick and yet with a reflectance of 0.82 and absorptance of 0.135 (compare to figure 10.4). This would give the overall spacecraft a lightness number of 6.0 (or 13.6 just for the sail). Presumably additional supporting structures would be required in practice, here considered as part of the payload. This sail is driven by a variable power laser, which can deliver a peak power of 58 TW. This exceeds the rate of energy use by all humans on the Earth (for all purposes), which was about 20 TW on average in the year 2022.

The optical array used to transmit the laser power (at 1 μm in the laser rest frame) has an enormous diameter of 1000 km. One motivation for such large optics is to use the laser also to brake the sail at the target star (see section 11.6.4). The Airy disk of this diffraction pattern has a diameter of only 100 km at Proxima Cen. When all of the peak laser power is intercepted by the sail, the average intensity on the sail's surface is $72\,\mathrm{kW\,m^{-2}}$, which is only about 50 times the intensity of the Sun at the top of the Earth's atmosphere. When the sail is nearer to the laser, the diffraction spot is smaller and so the intensity would be much larger, heating the sail to unacceptable temperatures. For this reason, the laser power is lower in the earlier parts of the journey.

The laser is used to accelerate the sail to a velocity of $0.5\,c$ over a period of 5.5 years – that's a lot of laser energy – at which point the spacecraft is 1.4 ly from the Earth. A further 5.9 years cruising results in a total travel time to Proxima Cen as measured on an Earth clock of 11.4 years. This is laser sailing on a grand scale, to say the least.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

11.6.2 Low-mass payloads

More recently there has been a trend towards designing extremely lightweight sails in order to bring the laser power and optic sizes closer to the range of the plausible. An example of such a study undertaken as part of the Breakthrough Starshot initiative is by Parkin (2018). This uses a 4.1 m diameter sail with an areal mass density of $2.0 \times 10^{-4}\,\mathrm{kg\,m^{-2}}$, corresponding to a mass of 2.6 g. The payload is tiny, just 1 g. The sail is assumed to have a reflectance of 0.7 and an absorptance of just 10^{-8}, so the overall lightness number of the sailcraft is 3.9. It is propelled by a laser beam with a maximum power of 200 GW, although the power is initially lower and then increased as the spacecraft recedes in order to compensate for diffraction losses. The laser power is channelled through an array of telescopes with an effective diameter of 2.7 km. The peak intensity on the sail is $8.7\,\mathrm{GW\,m^{-2}}$, which is why such a small absorptance is necessary. This intensity is considerably higher than that considered in the high-mass payload design above, which used used two orders of magnitude more laser power, but eight orders of magnitude more sail area. Despite this enormous intensity, the pressure on the sail is only 41 Pa (equation 10.19), yielding a total force of 540 N, yet the sail achieves an acceleration of $1.5 \times 10^{5}\,\mathrm{m\,s^{-2}}$, or 15 000 times standard gravity g. A bullet in a rifle is accelerated by a similar amount, but only for a few milliseconds. The sailcraft must tolerate this acceleration for several minutes.

Figure 11.17 shows how various properties of the laser and sail vary with time (measured in the laser rest frame). Initially, the diffraction pattern on the sail is smaller than the sail, and so the laser is not run at full power to avoid overheating. The acceleration profile is designed so that the sail temperature never exceeds 625 K. As the sail recedes, the diffraction spot gets larger, and the output power is increased. After three minutes, the angular size of the sail is too small to fill most of the beam and the number of photons lost around the sides becomes significant. After five minutes the maximum laser power is attained. The laser remains on for a total of eight minutes in the laser rest frame, but this corresponds to the sail receiving photons and being accelerated for nine minutes (also in the laser rest frame) due to the difference between coordinate time and retarded time (see figure 11.4). During this acceleration period the sail has travelled 67 light seconds, or 0.13 au, and has acquired its final velocity of $0.2\,c$. This corresponds to a kinetic energy (relative to the laser rest frame) of 1.9 GWh. As the laser has provided a total of 63 GWh of energy, the overall energy efficiency is just 3%. This is much lower than the value given by equation 11.9 (18%), because much of the energy never makes it to the sail, due in particular to diffraction losses around the side of the sail, but also absorption by the Earth's atmosphere.

11.6.3 Intermediate-mass payloads

Various laser sail designs were put forward in response to the Project Dragonfly competition workshop, organized by the Initiative for Interstellar Studies (i4is) in 2014. This required a sail propelled by a laser with 100 GW power or less to reach the Alpha Cen system within a century.

One solution was put forward by Häfner et al. (2019). They consider a sail made of a graphene monolayer of volume density $2210\,\mathrm{kg\,m^{-3}}$ with a thickness of 0.5 nm, and thus an areal density of $1.1 \times 10^{-6}\,\mathrm{kg\,m^{-2}}$. Adopting a reflectance of 0.05 and absorptance of 0.40, this corresponds to a lightness number of 345 for the sail alone. A sail 29.4 km in

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

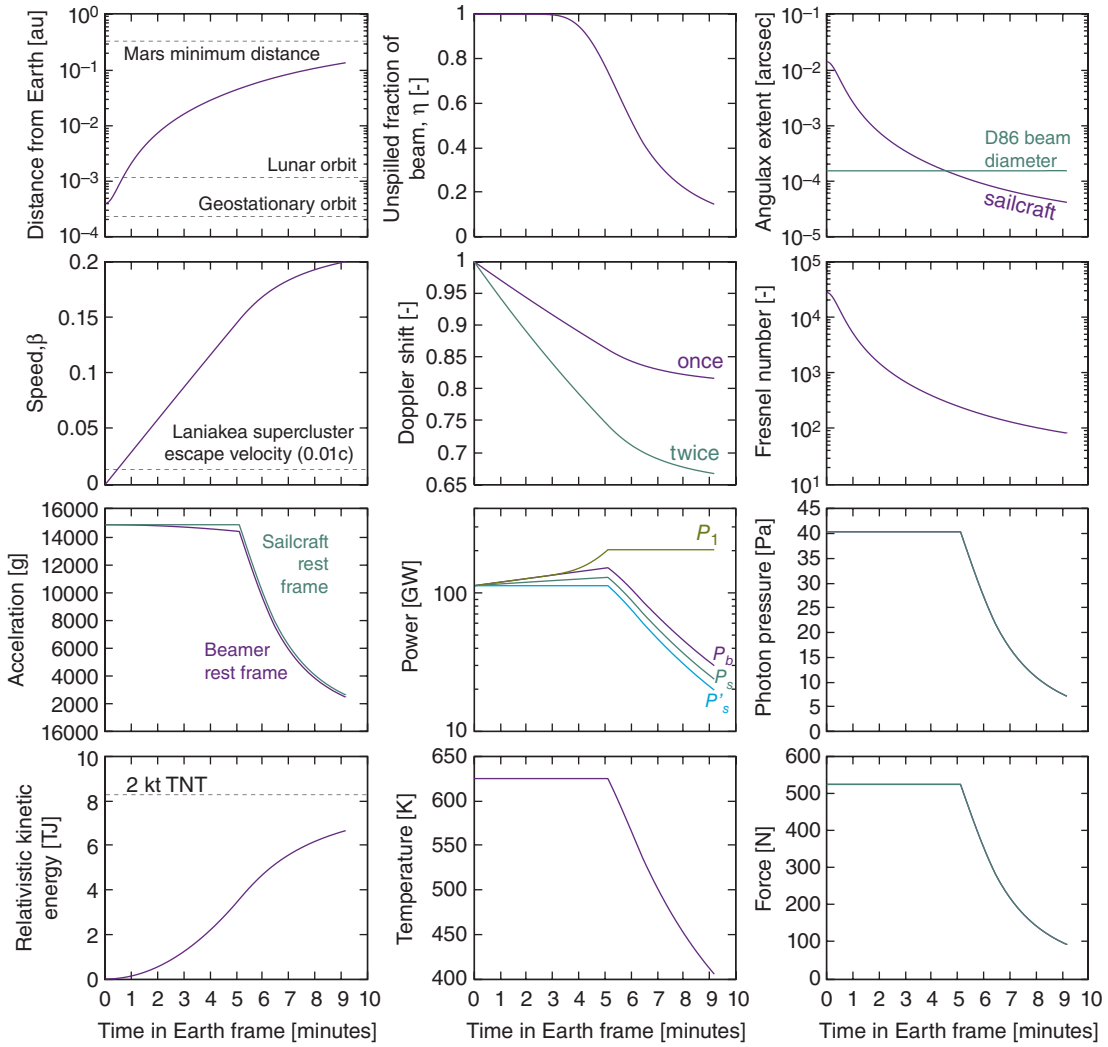


Fig. 11.17

The variation of various properties of a laser sail system as a function of time during the acceleration of a sail. In the power panel, P_1 is the power emitted by the laser in its rest frame, P_b is the power received by the sail in the laser rest frame (which is lower than P_1 due to diffraction losses), and P'_s is the power received by the sail in the sail rest frame (due to the double Doppler shift, equation 11.27). The timescale for P_1 is not correctly represented here as it does not take into account the photon travel time from laser to sail. Credit: From Parkin (2018), used with permission of Elsevier (conveyed through Copyright Clearance Center, Inc.) and the author.

diameter has a mass of 750 kg, and they adopt a total spacecraft mass of 2750 kg, of which 1000 kg is the payload and another 1000 kg is a magnetic sail used to brake the spacecraft at its destination, Alpha Cen. (We examine magnetic sails in chapter 13.) The overall lightness number is therefore 95.

A 100 GW laser with an effective aperture of 29.4 km (equal to the sail size) accelerates the sail at a nominal 0.061 m s^{-2} (equation 11.17) to 0.05 c over 7.8 years, after which the spacecraft is 12 000 au (0.19 ly) from the Earth. At this distance the sail has an angular radius of just 0.0082 nrad, or 1.7 μas. They use a 400 nm laser, so its diffraction angle

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

is 0.017 nrad. The laser has to be pointed with a higher accuracy than this, which is very challenging.[6]

At its cruising speed of $0.05\,c$ the spacecraft will arrive at Alpha Cen A after another 85 years. However, when the spacecraft is 7500 au (0.12 ly) from its destination, it jettisons the laser sail, powers up a current through its magnetic sail, and uses this to brake, first using the ions in the interstellar medium then using the higher density ions in the heliosphere (from the stellar wind) of Alpha Cen A. Adopting a very large current density of $J_s = 1.5 \times 10^{11}\,\mathrm{A\,m^{-2}}$, they compute that the magnetic sail can decelerate the spacecraft to a velocity of just $250\,\mathrm{km\,s^{-1}}$ by the time it reaches the star, taking around 21 years.

Dragonfly is a much more massive spacecraft than the one considered by Parkin (2018) above, and although they use a similar laser power, Dragonfly must use it for nearly eight years as opposed to eight minutes, and correspondingly consumes far more energy. Although the Breakthrough Starshot approach uses extremely light spacecraft with limited payloads (and essentially no braking ability), it could launch many more spacecraft within the same time period. Dragonfly, in contrast, permits a larger payload, and if it could go into orbit about the target star (or potentially a planet around it), it could perform significantly more science, potentially even dropping a small lander. The larger mass would also permit the spacecraft to carry more onboard power and carry larger communication systems for a higher capacity data link to the Earth.

11.6.4 Using a light sail to brake at the target system

By leaving the means of propulsion behind, a sailcraft overcomes the restrictions of the rocket equation. However, lacking an internal propulsion system limits the options for braking the spacecraft at the target system, whether to prolong the flyby or to go into orbit. We saw in the previous section that a massive sailcraft could be equipped with a magnetic sail to decelerate, but this is not compatible with a low-mass sailcraft, which could not carry the heavy equipment needed to generate a magnetic field.

A sailcraft of any size could use the light from the target star to decelerate. We can compute the change in velocity using equation 10.13, where v_∞ is the velocity of the sailcraft far from the star, and v_i is its velocity once it has approached to a distance r_i from the star. A 4 m diameter sail with a mass of 1 g and unit reflectance has a lightness number at the Sun of 19. When inbound to a Sun-like star with $v_\infty = 0.1\,c$, then by the time it reaches $r_i = 1$ au, its velocity has barely decreased: by only $0.5\,\mathrm{km\,s^{-1}}$. This might seem odd, given that this sailcraft starting at rest 1 au from the Sun would be accelerated to $v_\infty = 180\,\mathrm{km\,s^{-1}}$ (equation 10.13 with $v_i = 0$). But this is just a consequence of the quadratic dependence of kinetic energy on velocity: subtracting $180\,\mathrm{km\,s^{-1}}$ in quadrature from $0.1\,c$ makes hardly any difference.

We can only start to achieve non-negligible decelerations from the star if the sail (a) has a much higher lightness number, (b) has a smaller incoming velocity, and (c) approaches

[6] Häfner et al. (2019) adopt a larger acceleration of $0.127\,g$, which appears to be based on their erroneous equation 1, which has the term $1+\tilde{r}$ in place of $2\tilde{r}+\tilde{a}$ in equation 10.19. They therefore quote an acceleration distance and time that are about half the values I obtain using the constant acceleration formulae (speeds are barely relativistic, so we don't need to use equations 11.19 and 11.22). On the other hand, there will have been significant diffraction losses later in the journey, which have not been taken into account.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

much closer to the star. Adopting a lightness number of 410 (e.g. zero-reflectance graphene with absorptance of 0.4), an initial velocity of 0.05 c, and getting it to approach to five solar radii, then the spacecraft can be decelerated by 1080 km s^{-1} to 0.046 c, which is still not a significant decrease in speed.

The above calculations assume the target star is the Sun. If the star has a larger luminosity-to-mass ratio, then it will provide a larger deceleration. Alpha Cen A (section 2.5.2) has a luminosity-to-mass ratio about 1.4 times that of the Sun, and therefore sails have a lightness number at Alpha Cen A that are also larger by this factor (equation 10.10). Taking advantage of this, and also using an even lighter sail and a slightly lower initial velocity, it may be possible to decrease the sail velocity at Alpha Cen A to around half its initial value. Furthermore, the sail's path can then be deflected in a gravitational – actually, a photo-gravitational – assist toward Alpha Cen B, where it is decelerated again, and then redirected to Proxima Cen. Provided the incoming velocity is below some maximum value, it should be possible to put the spacecraft into a range of desired orbits around Proxima Cen.

An ingenious if somewhat impractical idea for braking the sailcraft using the laser itself was outlined by Forward (1984) and is summarized in figure 11.18. The sail has a diameter of 100 km and a mass of 785 t. The sailcraft is accelerated continuously over most of the distance to the target star (Alpha Cen). This reduces the power required compared to a shorter acceleration distance, but it also requires a large focusing array to avoid diffraction loses. For this purpose, a 1000 km Fresnel lens is put into space near to the Sun to focus the laser beam. Operating at 1 μm, this lens produces a diffraction pattern with 100 km diameter at the distance of Alpha Cen. The laser power is nonetheless enormous at 7.2 TW (a third of the world's power in 2022). After 40 years of illumination at a leisurely acceleration of

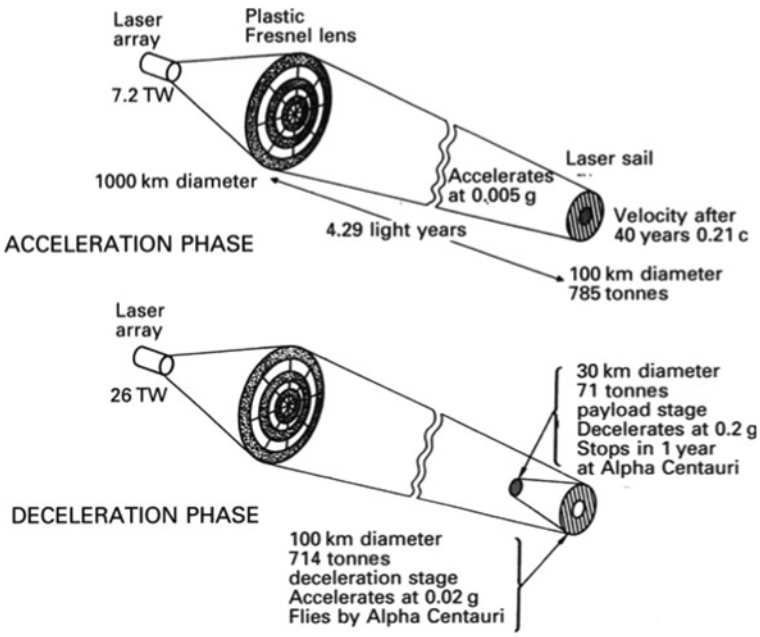


Fig. 11.18 A two-part laser sail that uses the Earth-based laser to decelerate at the destination. Credit: Used with permission of Springer Nature BV, from Turner (2009); permission conveyed through Copyright Clearance Center, Inc. Based on figures in Forward (1984).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

0.005 g, the sail reaches a velocity of 0.21 c and has travelled 4.29 ly, and so is nearly at its destination. At this point the sail splits into a central disc of diameter 30 km (71 t), and an outer annulus (714 t). The disc – which carries the payload – is then turned around so that its highly reflecting side faces the target star. Its non-reflecting side faces the laser, and as it has a higher areal density than the rest of the sail, its acceleration (from absorption of the laser light) is now much lower. The annulus continues to be accelerated by the laser as before and so moves away from the disc. The annulus now reflects laser light back onto the disc – it sees the reflecting side – thus decelerating the disc. The laser power is increased to 26 TW and after one year the disc has decelerated to zero velocity at Alpha Cen. The annulus continues past Alpha Cen and is lost.

This ambitious idea obviously has many challenges, not least the enormous size of the optics, the huge laser power, the need to run the laser continuously for over 40 years, and the heating of the payload. But it could even be configured to not only stop the sailcraft at the target star, but even bring it back to Earth.

11.7 Summary

Laser-propelled sails offer an alternative to nuclear rockets for high-speed interstellar travel. They are conceptually similar to solar sails, but operate at much larger radiation intensities. Laser power is transmitted by large telescope arrays on the Earth (or Moon or orbit). Even with a very large array (kilometres), diffraction of the beam results in significant amounts of light spilling around the edge of the sail once it exceeds some distance. Most of the laser light should therefore be delivered within a limited distance, and therefore time, which demands a very large laser power. To accelerate even an ultralight sail a few grammes in mass and a few metres in diameter to 0.1–0.2 c requires hundreds of gigawatts of power, ideally delivered within several minutes. The energy efficiency is of the order of a few per cent, with most of the energy lost either to diffraction spillage or unavoidably in the reflected photons (figure 11.3).

Due to the much higher intensities, laser sail materials face much larger thermal demands than solar sails. Metallic films are too absorptive with low melting temperatures. Dielectric materials – such as molybdenum disulphide and silicon nitride – are preferable. Surprisingly, low-reflectance, high-absorptance materials may perform well if they are light enough and have high emissivities and melting temperatures. One such example is graphene. Another approach to materials engineering is to design nanophotonic structures to achieve high reflectance and low mass.

As the laser array must be of the order of kilometres in size, a single filled aperture is implausible. It must instead be synthesized as a compact array of many smaller apertures (e.g. tens of thousands each of 10 m diameter). These need to deliver a coherent wavefront, so their phases must be carefully controlled. If placed on the surface of the Earth, adaptive optics is required to correct the effects of atmospheric turbulence that would otherwise break the coherence.

Numerous laser sail mission concepts have been developed, some more plausible than others. Those originating in the 1980s tended to involve massive sails and payloads, on the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012

scale of hundreds of tonnes and many kilometres in diameter. They required tens of terrawatts of power maintained for years, propelled by arrays hundreds of kilometres in size.

More recent designs have focused on ultralight sailcraft. Their low mass (grammes) severely limits the scientific instruments they can carry. However, each sail is accelerated only for a few minutes or hours. As the marginal cost of a sail is small compared to the expensive optical array, it makes sense to launch many sailcraft, perhaps hundreds or thousands over the course of a year. In this way, the scientific instruments and objectives can be spread over many spacecraft. We will look at this distributed science concept more in section 16.3.

11.8 Exercises

1. Show that equation 11.7 gives the velocity of the sail relative to its initial rest frame when $q = hv'_e/mc^2$.
2. Figure 11.3 shows that the overall efficiency of energy transfer from the laser to the sail increases with sail velocity β. To reach a large velocity where the efficiency is high, the sail had to pass through a series of lower velocities where the overall efficiencies were all lower. How can we reconcile this?
3. Show that equations 11.20 and 11.23 follow directly from equation 10.19 in the non-relativistic limit.
4. In section 11.1.5 we used the Doppler expression to compute the rate of arrival of photons at the sail in the sail rest frame $\dot{n}'_s$, in terms of the rate of emission of photons from the laser in the laser rest frame $\dot{n}_l$. Derive this instead in two steps, first using equation 11.11, which is purely classical, and then using the time dilation (to relate $\dot{n}_s$ to $\dot{n}'_s$). Plot the rate ratios vs β for these two steps, as well as the rate ratio in the final expression (i.e. $\dot{n}'_s/\dot{n}_l$)
5. The diffraction angle of the laser beam observed on the moving sail is larger than the diffraction angle at the laser source due to aberration, and can be described by equation 9.5. Show that this expression can also be obtained from equation 11.30 using the Lorentz transformations and/or Doppler effect.
6. Modify the derivation of the equation of motion in section 11.1.3 to accommodate the diffraction losses discussed in section 11.2.1. To keep it analytic, assume all the power is distributed uniformly over the Airy disk.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 11:25:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009700429.012