Subsystem Design: I Like Your Attitude

8

IN the last chapter we discussed the overall process of spacecraft design, and showed where the requirements for the design of the spacecraft subsystems come from (see Table 7.1 and Figure 7.1). The subsystems are there simply to support the payload in achieving the spacecraft's mission objective, so that the resources they need to do this job governs the way they are designed. In this chapter and the next, we take a closer look at the design of the major spacecraft subsystems to identify the main drivers that dictate their design. (The mission analysis subsystem was discussed in Chapters 2, 3, and 4.) Discussing the design of each subsystem will help explain why spacecraft look the way they do.

This chapter discusses the attitude control subsystem (ACS), which is perhaps the most complex of the subsystems and is very influential in determining the overall shape of the spacecraft. *Attitude control* may sound like a sinister governmental plot out of George Orwell's novel *1984*, but in the context of spacecraft engineering it refers to controlling the rotation of the vehicle.

Attitude Control Subsystem

What Does the ACS Do?

Table 7.1 in Chapter 7 stated the purpose of the ACS: "to achieve the spacecraft's pointing mission." What does this mean? Most operational spacecraft in orbit have payloads that require pointing. For example,

- a communication satellite (comsat) needs to point its communications dish(es) at a ground station to receive and transmit the stream of telephone conversations that are using the system,
- an Earth observation spacecraft needs to point its imaging cameras at the targets of interest on the ground, and
- a space observatory needs to respond to ground commands to point a telescope at particular objects (planets or galaxies) in the sky.

G. Swinerd, *How Spacecraft Fly: Spaceflight Without Formulae,*
DOI: 10.1007/978-0-387-76572-3_8,

So, the ACS addresses the pointing, or rotation, of the spacecraft. In earlier chapters we talked a lot about orbits, focusing on how the spacecraft's center moved along its orbit. By contrast, in this section about the ACS, we are not concerned with the motion of the spacecraft's center along its orbit, but rather with the way the spacecraft rotates *about* its center. We can imagine ourselves being in the same orbit as the spacecraft, just a few meters away, and watching it rotate in response to commands to point the payload instruments. The word *attitude* describes the position of the spacecraft in a rotational sense, and a *change in attitude* therefore implies a rotation. What causes the spacecraft to rotate (or to stop rotating)? In the earlier chapters on orbits, we saw that it was *force* that changed the orbit along which the spacecraft moves. In contrast, it is *torque* that causes the spacecraft to rotate about its center.

We discussed torques in Chapter 3, when we looked at gravity anomaly orbit perturbations. A torque is a rotational force such as the one we apply to remove a bolt from a wheel when we change a flat tire. We apply a force of a number of Newtons in a rotational sense by pushing down on the end of the handle of the wrench. The size of the applied torque is based not only on the amount of force exerted but also on the length of the handle of the wrench. The longer the handle, the greater the *moment arm* and the more torque there is. It is easier to loosen the wheel bolts if the handle of the wrench is lengthened. Effectively, this increases the amount of torque, without requiring us to apply more force on the handle. The magnitude of the torque is given by the force times the moment arm, and has units of Nm (Newton meters).

The main job of the ACS is to control the rotational state of the spacecraft by using on-board devices called *control torquers* that produce torques on command to rotate the spacecraft. The most obvious kind of control torquer is a pair of thrusters fired in such a way as to produce a rotation. We will discuss thrusters later, when we talk about the propulsion subsystem, but essentially they are small rocket engines (small enough to be held in the palm of your hand), each producing a thrust force on the order of a few Newtons (a force of 1 Newton is about the weight of a small apple). They are located in clusters around the exterior of the spacecraft, and can be fired in opposed pairs (see Figure 8.4) to produce a torque, and therefore a rotation of the spacecraft. Other kinds of control torquers are discussed later. However, in the interim, it is helpful to briefly describe the main functions of the ACS, as these are the aspects that lead to the ACS design:

- To achieve the pointing mission of the payload, in terms of directions and accuracy. A comsat, for example, may need to point its antenna at a ground station with an accuracy of 0.1 of a degree, whereas a space observatory such as the Hubble Space Telescope may need to point the telescope at a galaxy with an accuracy of less than an *arc second*. If we divide a degree by 60, we get an arc minute, and if we divide an arc minute by 60 we get an arc second. So an arc second is a tiny angle, being 1/3600th of a degree!

- To achieve the pointing requirements of other subsystems, a process sometimes called *housekeeping*. For example:
  - Pointing a solar panel at the Sun to generate electrical power
  - Pointing an antenna at a ground station on Earth's surface to downlink payload data
  - Pointing thermal radiators to the cold of space to allow heat to escape (this will make more sense after we look at the thermal control subsystem)
  - Pointing a rocket engine in the right direction before it is fired, so that the correct change in the orbit is achieved
- Overall, to manage the rotational state of the spacecraft, meaning motion and torques *about* the spacecraft's center.

How Does the ACS Work?

Another aspect that influences the ACS design is how the ACS operates to achieve these functions. Figure 8.1 shows a typical ACS operation, and introduces the main hardware components that comprise the ACS. Starting at the top of the figure, torques act on the spacecraft and cause it to rotate. There are two types of torque. First, there are those we apply deliberately to control the spacecraft's rotation using the on-board ACS hardware. These are produced by the control torquers that we mentioned above. Second, there are *disturbance torques*, which are the rotational equivalent of orbit perturbations. As we saw earlier, the orbit motion of the spacecraft was altered by perturbing forces (see Chapter 3). The spacecraft is also disturbed in terms of its rotation by naturally occurring torques, produced by the spacecraft's interaction with its environment. For example, Figure 8.2 shows how a disturbance torque is generated by a spacecraft's interaction with the atmosphere. As we saw in Chapter 3, atmospheric drag causes a decrease in orbit height, but it can also produce a disturbance torque that causes an unwanted rotation in the spacecraft, which needs to be corrected by the ACS.

The torques act on the spacecraft and cause its attitude to change, that is, cause it to rotate. Moving around Figure 8.1, we see that this rotation is detected and measured by *attitude sensors*, which are hardware components of the ACS, usually optical sensors looking out at *reference objects*, such as the Sun or stars. To picture how these work, imagine being inside a box with windows, which is good description of an airplane. Imagine sitting in a

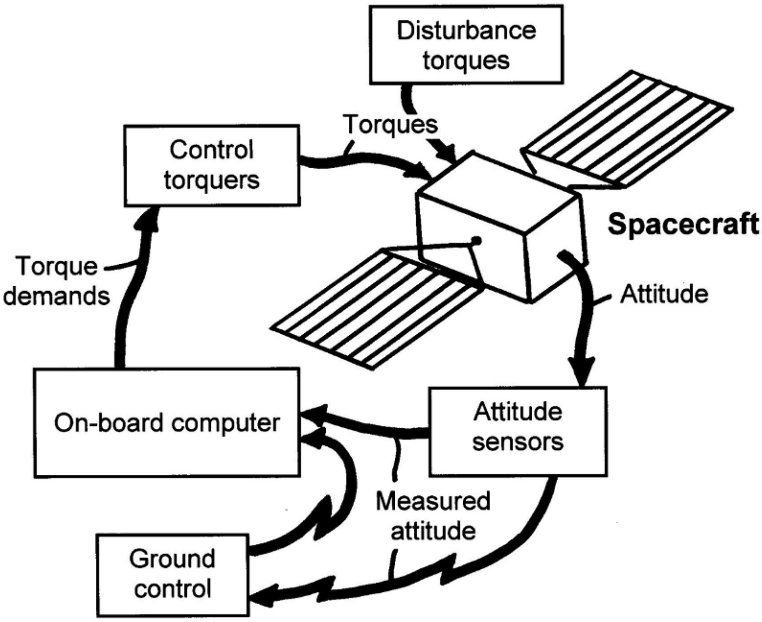

Figure 8.1: A block diagram showing how the ACS operates.

window seat on a night flight. If the cabin lights are down, the stars can be seen easily, and while the airplane is not turning (rotating), the stars look as if they are stationary in the window. However, as soon as the aircraft starts to turn, they appear to move across the window. In the same way, the sensors look out of the spacecraft and interpret movement of reference objects, such as the stars, as a rotation of the vehicle—so allowing measurement of the rotation.

The sensor measurements of the spacecraft's rotation are passed to the on-board computer, which itself can be considered to be another piece of ACS hardware. The measurements are processed by *control software*, which is essentially fancy mathematics coded into the computer to calculate the spacecraft's attitude. This estimate of the attitude is then compared with the attitude required to achieve the pointing mission, and if they differ, the computer's control software calculates what torques are required to correct the spacecraft's attitude. These *torque demands* are then passed to the control torquers (continuing our walk around Fig. 8.1), which then apply the required torques to bring the spacecraft's attitude back to where it should be.

For example, the *pointing mission* for the spacecraft may be to point the antenna of a communications satellite to a ground station to facilitate intercontinental telephone calls. If the satellite is disturbed, the pointing of

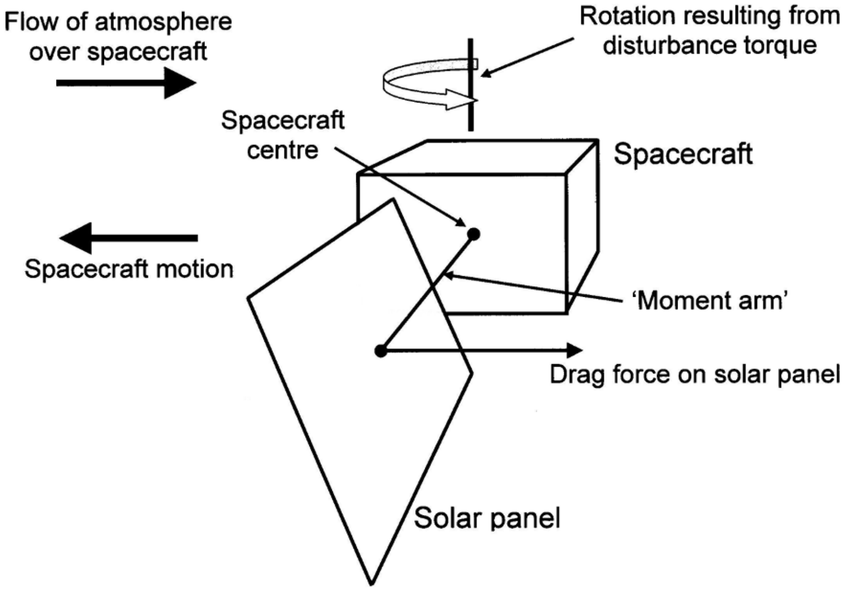

Figure 8.2: An example of a disturbance torque—in this case that produced by aerodynamic drag.

the antenna to the ground station will be affected, but this disturbance will be detected by the ACS sensors. The computer will then process the sensor measurements to command the on-board torquers to correct the pointing error, thus maintaining the communications link. One thing to note about the ACS operation shown in Figure 8.1 is that it operates in a loop, and on most spacecraft it does this automatically many times a second, so that the pointing mission is continuously monitored and maintained. This looping-type operation is referred to as a *feedback loop* by the ACS engineer.

At the bottom of Figure 8.1 there is a ground intervention in this automated process. For example, to operate a space telescope, the astronomers on the ground need to command the telescope to point at a particular galaxy, for example. They can do this by typing the position of the object of interest into a computer on the ground, and this information is then up-linked to the spacecraft and processed by the on-board computer to produce torque demands, which are then executed by the control torquers to rotate the spacecraft to direct the telescope to the required segment of sky.

The main functions of the ACS help other subsystems in their operation (e.g., pointing a solar panel to the Sun to help the power subsystem to do its job). Similarly, if we look at the typical ACS operation outlined in Figure 8.1, we can see that other subsystems help the ACS do its job (e.g., sensors, computers, and control torques need electricity from the power subsystem to work). So, in the design process, the ACS engineer has to work together with many other subsystem engineers, resulting in a very interactive design process.

Attitude Stabilization

One reason why the ACS is considered to be such an important element of the spacecraft is that the type of attitude stabilization used on a particular spacecraft is very influential in determining what it looks like. There are four general types of stabilization, which are illustrated in Figure 8.3. Types 1, 2, and 3 involve spinning all, or a part, of the spacecraft. This spin feature makes the spacecraft's attitude inherently stable; if the spacecraft is affected by a disturbance torque, the change in attitude that results is small. This is a useful feature, as it means the ACS does not have to work so hard to control the spacecraft's attitude.

We can get a good idea of how this inherent stability due to rotation works by looking at a bicycle. The bicycle's tires provide two small points of contact with the ground, each perhaps a couple of centimeters across. The bicycle rider represents a large mass on the top. Other objects with these two characteristics tend to fall over; for example, no matter how hard we try, we cannot balance a nail on its point. But strangely the bike rider is quite happy whizzing along the road without any thought of the bike toppling over. Why? It's basically because the rotation of the wheels give the bike stability; the axes about which the wheels rotate become stiff, in the sense that they want to remain pointing in the same direction. The wheel axles remain horizontal, ensuring that the bike stays upright. As the rider slows down, the wheels' spin rate correspondingly decreases. Eventually the wheels stop rotating, and then the bicycle's stability is lost, and the rider has to put a foot down on the ground to prevent the bike from toppling.

And so it is with a spacecraft. When it spins, the spin axis becomes stiff, tending to make the spacecraft point in a fixed direction. The spin axis becomes less sensitive to disturbances, giving the spacecraft as a whole this characteristic of inherent stability. A spacecraft with this type of spin stability is said to have *momentum bias.* In Figure 8.3, a spacecraft with type 1 stabilization is called a *pure spinner.* These are usually cylindrical in shape, and the whole spacecraft rotates at a rate of a few tens of revolutions per minute (rpm), providing spin stability. The example shown is the Meteosat SG (second generation) satellite, which is a spacecraft that provides satellite pictures for weather forecasts. With this type of stabilization, there is no part of the vehicle or its contents that is not rotating; thus, it is difficult to accommodate payloads that need to point in a fixed direction.

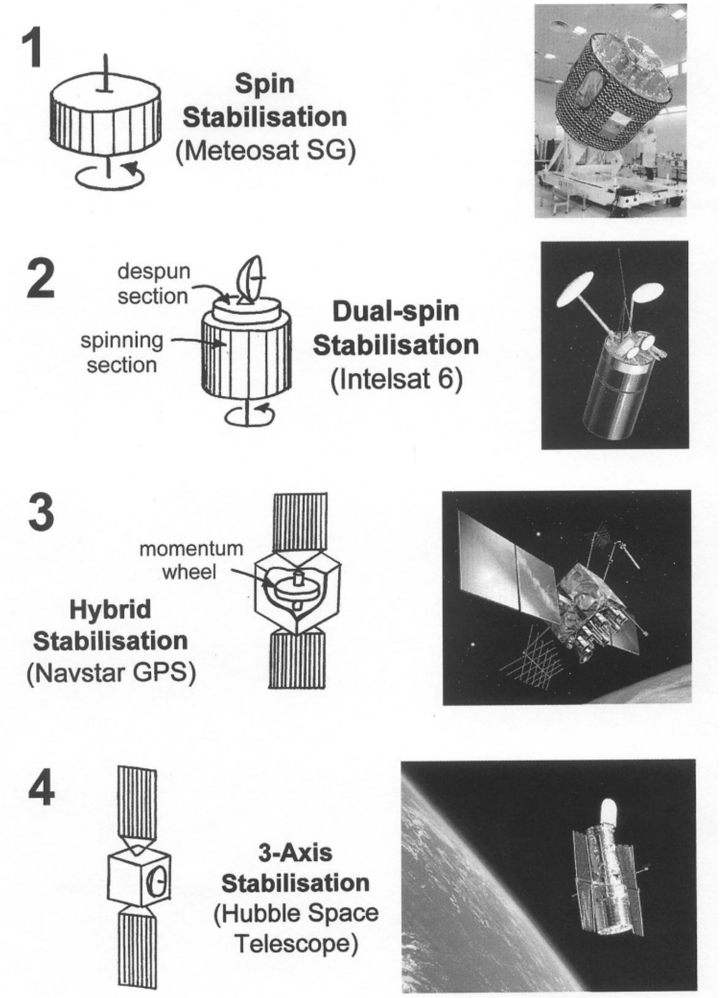

Figure 8.3: The four general types of spacecraft attitude stabilization. All spacecraft fall into one of these categories, some examples of which are illustrated. (Image credits: Meteosat SG image, copyright © ESA; Intelsat 6 image, copyright © Boeing; GPS Navstar 2R image, copyright © Lockheed Martin; Hubble Space Telescope image, copyright © NASA.)

This problem is alleviated by the use of type 2 stabilization. Spacecraft with this stabilization type are referred to as *dual spinners*, and have a cylindrical section spinning at a few tens of rpm (like type 1), giving it spin stability. However, there is a platform mounted on top of the vehicle that is mechanically de-spun, where payloads that need to be pointed in a fixed direction, such as antennas or imaging cameras, can be mounted. The example shown in the figure is the Intelsat 6 communication satellite, which remains (at the time of this writing) the largest example of this type of attitude stabilization.

The third type, which I have called *hybrid stabilization* for want of a suitable label (although this terminology is not used universally by ACS engineers), is quite an interesting arrangement. Here the spacecraft acquires spin stability by mounting a momentum wheel inside the vehicle (see Figure 8.6). In this case, the spin stability is achieved by spinning a small object very rapidly (the wheel rotates at a typical rate of a few thousand rpm), rather than spinning a big object more slowly (the whole or part of the spacecraft rotating at a few tens of rpm, as is the case for the pure and dual-spinners). The mass of the wheel is typically a few kilograms, and its spin rate is maintained at around 6000 rpm. Given that the wheel is rigidly mounted in the spacecraft, its spin stability is transferred to the spacecraft as a whole. Thus the vehicle has the benefit of inherent stability, while at the same time allowing lots of space on the exterior surfaces of the spacecraft to mount payloads and deploy solar panels. The example shown in Figure 8.3 is that of a Navstar GPS satellite, which is part of the U.S. Department of Defense's constellation of satellites used for navigation. Like all such spacecraft using this type of stabilization, the spin stability is not at all obvious to the casual observer, as the mechanism for achieving this (the momentum wheel) is concealed inside the vehicle.

Type 4 is referred to as *three-axis stabilization*. In this case the spacecraft has no significant rotating parts, and thus does not have the inherent stability associated with the other types. As a consequence, the ACS has to work a bit harder to achieve the pointing mission. This lack of stability seems on the face of it to be a disadvantage, but often it is the only suitable option. A good example of this is a space observatory, such as the Hubble Space Telescope shown in Figure 8.3. To achieve its pointing mission, it has to rotate freely to point in various directions, and as a consequence there is no axis in the spacecraft about which it is sensible to employ spin stability. This would make the spin axis stiff, which would make no sense at all if it has to be moved around a lot in the process of pointing the telescope.

The different types of stabilization affect the overall shape or configuration of the spacecraft, which is why the control engineer would say that the ACS is the heart of the spacecraft.

Control Torquers

We have already briefly mentioned control torquers in our walk around the control loop in Figure 8.1. These are items of ACS hardware that are essentially the muscles of the ACS, converting the virtual commands produced by the on-board computer into physical torques that rotate the spacecraft.

Thrusters

Perhaps the most obvious way of producing a control torque is to fire two thrusters in opposite directions, as illustrated in Figure 8.4a. These small rocket engines are set up in groups, called *thruster clusters*, which are located at a number of positions around the spacecraft to ensure that the attitude and orbit control functions can be effectively achieved. By firing thrusters in pairs using different clusters, we can rotate the spacecraft in any direction as shown in Figure 8.4b.

Magnetorquers

When I was a child, I made an electromagnet using a large 6-inch nail, some wire, and a battery. I wound the wire many times around the nail, and

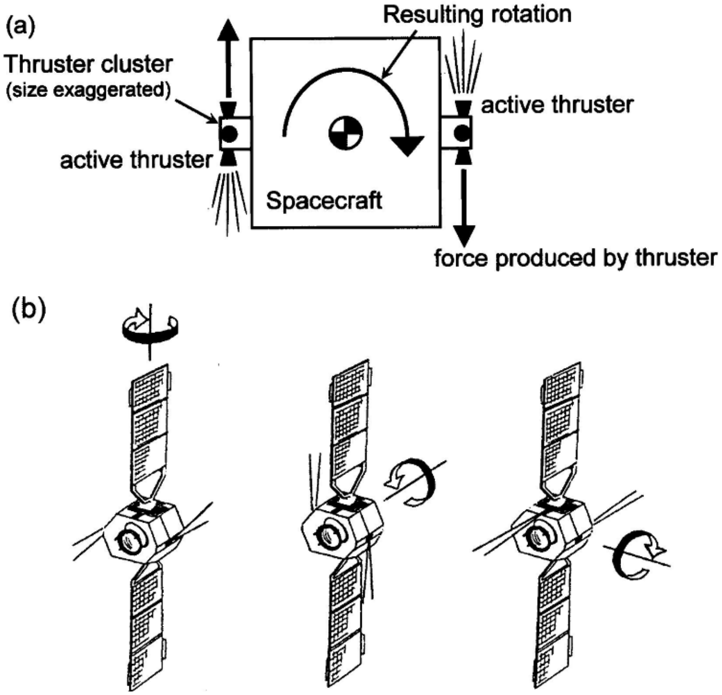

Figure 8.4: (a) The firing of two thrusters in opposed pairs produces a rotational force (torque), causing the spacecraft to rotate. (b) Firing opposed pairs in appropriately chosen thruster clusters allows the spacecraft to be rotated in any direction in three dimensions.

attached the ends of the wire to the battery, and magically when the circuit was made the nail became a magnet. I remember being intrigued by this, and enjoyed the childish pleasure of picking up paper clips and toy cars with this magnet, which I could turn on and off simply by making or breaking the contact with the battery. This simple homemade electromagnet is the basis for another form of control torquer, the *magnetorquer rod*, although the real thing is a little more precisely engineered and a bit bigger! The nail is replaced by a metal rod, usually made from an alloy containing iron, and it can range in length from half a meter to about 2 meters, depending on the size of the spacecraft that needs to be torqued. A considerable length of wire is then wound around this core, giving us an electromagnet that can be activated on command by passing an electrical current through the device.

How is this used to rotate a spacecraft? Consider a compass; the compass needle is simply a magnet mounted on a pivot to allow it freedom to move. It points north because, as a magnet, it tries to align itself with the local magnetic field lines, which at Earth's surface run south to north (see Figure 6.3a in Chapter 6). In the same way, if a current is passed through a magnetorquer, it becomes a magnet, and as a consequence it too will tend to rotate to align itself with the local magnetic field at its orbital position, as shown in Figure 8.5a. Since the magnetorquer rods are firmly attached to the spacecraft (Fig. 8.5b), the vehicle will also share this rotation. So, if we know where we are in orbit, and what the magnetic field is like there, we can produce control torques to rotate the spacecraft on command, by simply passing electrical current through the appropriate magnetorquer. The idea sounds simple enough, but the implementation of this type of control is a fairly complicated business for the ACS engineer. Despite this, however, magnetorquers are commonly used on spacecraft. For example, the Hubble Space Telescope (HST) is equipped with magnetorquer rods to generate torques on the vehicle. An advantage of using magnetorquers on spacecraft like the HST is that they are clean, unlike thrusters that squirt out propellant each time they are used, which could end up contaminating the sensitive telescope optics. Figure 8.5c shows a couple of 1-m magnetorquer rods mounted on a test bed prior to installation in a spacecraft.

Reaction Wheels

Another commonly used type of control torquer is the *reaction wheel*. These are precisely engineered wheels, usually about 15 to 30 cm (6 to 12 inches) in diameter, with a mass on the order of a few kilograms. Their actual dimensions are governed by the size of the spacecraft in which they are installed, and how fast the spacecraft needs to rotate to achieve its pointing mission. To rotate the spacecraft in any direction in three dimensions, three

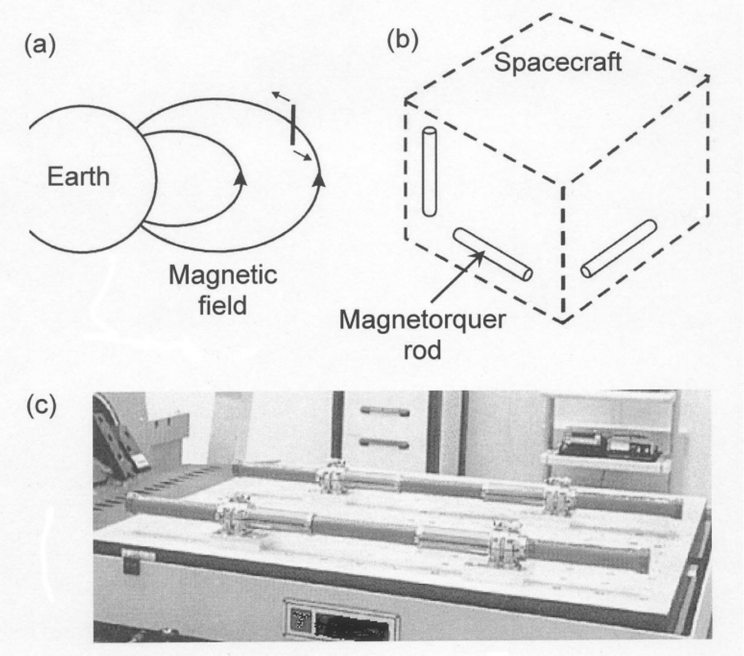

Figure 8.5: (a) When turned on, the magnetorquer rods align themselves with the local magnetic field. (b) The magnetorquer rods are firmly attached to the spacecraft, so that their rotation is shared by the vehicle as a whole. (c) Two magnetorquers under test. (Image courtesy of Dutch Space.)

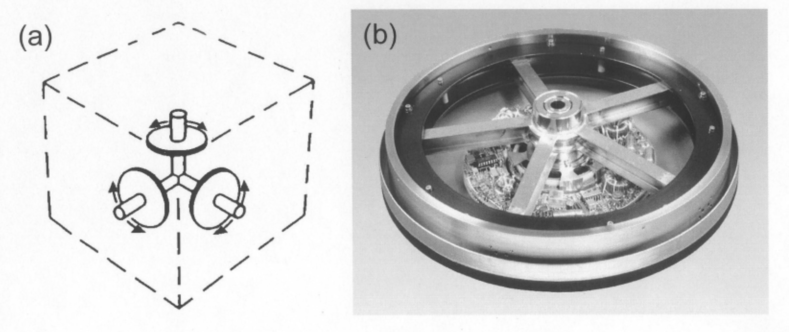

Figure 8.6: (a) Three reaction wheels are mounted rigidly in the spacecraft with their axes perpendicular to one another, to allow rotation of the spacecraft in any direction in three dimensions. (b) A typical reaction wheel. The wheel is usually sealed inside a disk-shaped canister, the lid of which has been removed in the picture. The electronics that control the wheel can be seen underneath the wheel in the base of the canister. (Image courtesy of Rockwell Collins Deutschland.)

wheels are usually employed with their spin axes mounted at right angles to each other, as shown in Figure 8.6a. However, the ACS engineer will usually mount a fourth wheel for redundancy reasons, with its spin axis canted at an angle to the other three, to allow control of the spacecraft in the event of the failure of one of the three primary wheels. An example of what a reaction wheel looks like is shown in Figure 8.6b.

To see how reaction wheels work to rotate the spacecraft, let's focus on just one of them. The wheel is connected to a torque motor, which itself is rigidly attached to the structure of the spacecraft. A torque motor is simply an electric motor that can be used to spin the wheel, a bit like the motor found in a domestic power drill; when we squeeze the drill's trigger, an electric current passes through the torque motor in the drill to rotate the business end—and so it is with the reaction wheel. To rotate the spacecraft around the axis of the wheel, electric current is passed through the wheel's torque motor, and as a consequence the wheel spins. To understand how this causes the spacecraft to rotate, we return to the power drill. If we "tweak" the drill trigger, the chuck and drill bit will spin in one direction, but the handle of the drill kicks back (in a rotational sense) in the opposite direction. This is why astronauts have trouble using power tools when working during space walks; the reaction causes them to rotate as well as the tool, so they have to be firmly attached to the spacecraft to prevent a rather undignified pirouette! It is this same reaction that kicks the spacecraft into rotational motion about the wheel axis, but in the opposite direction to the wheel's rotation. To summarize: To rotate the spacecraft about the wheel's axis, an electric current is applied to the wheel's torque motor. The wheel spins, and as it does so it produces a rotational kick in the opposite direction on the torque motor. Since the torque motor is attached to the spacecraft's structure, this kick is transferred to the spacecraft, which in turn begins to rotate about the wheel axis. This is an application of Newton's third law of motion, which we talked about in Chapter 1: for every action there is an equal and opposite reaction.

But how can the spacecraft be brought to rest again, because once set rotating, the spacecraft will continue to spin forever, as there are no frictional or other forces in space that will stop it. To stop the rotation of the spacecraft, the wheel is brought to rest. The braking (slowing) of the wheel produces a reaction on the spacecraft in the opposite direction that will slow and stop the rotation.

This method of changing a spacecraft's attitude is clean, efficient—elegant even—and just requires a bit of electrical power, which is (usually) freely available through the conversion of sunlight by solar panels. So it doesn't cost propellant, as is the case for thrusters, and generally it has a larger torque capability than magnetorquers. It also works equally well in finely pointing a spacecraft in a particular direction as it does in rotating the vehicle through large angles. Thus, it is the most commonly used technique by the ACS engineer to control the rotation of the spacecraft.

Summary

The choice of the attitude stabilization type has considerable impact on the overall shape and appearance of the spacecraft, and the ACS design is highly interrelated with the design of other subsystems on board the spacecraft. The ACS helps other subsystems in their operation, and conversely the ACS requires their services to allow its own operation.

In terms of hardware, the ACS is composed of sensors, control torquers, and computer processors, but one thing we have not discussed is the complexity of the mathematical control algorithms that are programmed into the onboard computer to make it all work. As a result, the subsystem engineer who works in this area needs to be not only good at the hardware design and its integration, but also a mathematician.

In the next chapter, we discuss the other subsystems and their design.