3 Spacecraft Attitude Control

3.1 A Distant Rocking

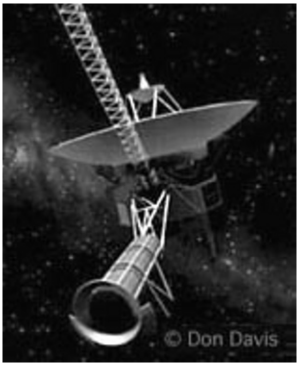

Fig. 3.1. One of the twin *Voyagers*, viewed from the end of its power supply boom. The science boom, not visible here, extends out the opposite side. Adapted from animation cell © Don Davis, reproduced by permission.

Today *Voyager 1* is sending telemetry via Station 14 in California's Mojave Desert, one of the Deep Space Network's 70-meter diameter tracking antennas, to the handful of engineers responsible for the craft.

Among the computer screens full of numbers and plots here in *Voyager*'s realtime operations support area, we take notice of a single data display that shows the spacecraft tracing its constant, slight changes in attitude, its orientation in space, rocking back and forth hour after hour. The craft is never still. As many spacecraft do, it is slowly oscillating about its three axes under computer control. The graph before us clearly shows three lines of measurements gradually building themselves into a plot, as radio symbols finish crossing the 16.2-trillion-meter distance,[1] a journey of fifteen hours at the speed of light. In seconds, the symbols are decoded in the desert, error-free, into bits of telemetry data. Milliseconds later, a program running on the computer in front of us is parsing those bits, and displaying some of them on this screen as points on the graph.

There is a vertical scale to the left of each line whose values range from zero at center, up to +0.10° and down to –0.10° (see Figure 3.2). Measurements in three *telemetry channels,* which we discussed in Chapter 1, appear on this plot. Labeled "PITCH, YAW, and ROLL," they keep reporting on *Voyager*'s relentless attitude changes, extending their individual lines pixel by pixel. A few pixels appear every minute, building from left to right on the screen. After a few hours, the three traces of data points, including the sawtooth-shaped yaw trace, reach the right-hand side of the screen, and the display scrolls the graph over to make more room to continue plotting. Every bit of this, and other engineering data from *Voyager*,

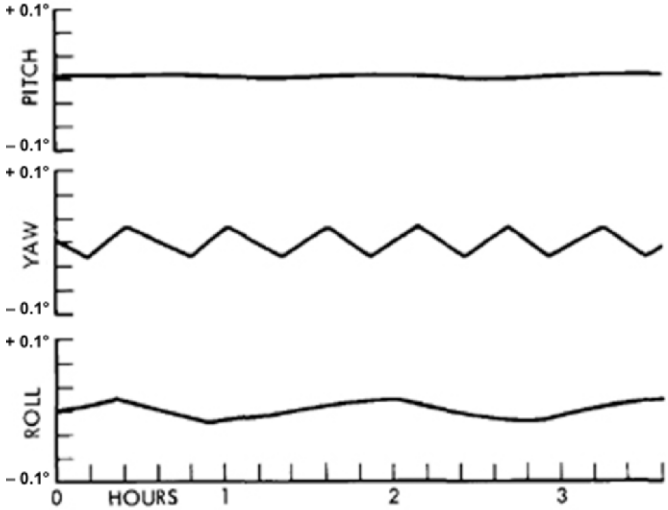

Fig. 3.2. *Voyager 1*, over sixteen thousand million kilometers distant, has been reporting its constant changes in pitch, yaw, and roll for more than thirty-one years. As usual, there happens to be hardly any pitch motion. All the elbows in the yaw line, and the first two in the roll line, show when the attitude control system caused a thruster to fire. Deadbands other than the 0.1° values shown here may be selected via command. Courtesy NASA/JPL-Caltech.

is being stored for analysis and maintained in an off-line archive. And of course the science data from *Voyager 1*'s six functioning instruments, as they sense the environmental conditions out past the solar wind's termination shock, within the *heliosheath*, is being stored and distributed to *Voyager* Project Scientist Ed Stone (1936–) and his teams of investigators.

The total mass of the spacecraft including propellant is about 730 kilograms as of late 2008. Each thruster firing exerts a little less than 0.9 N of force at an arm of about half a meter from the spacecraft's center of mass. To get an idea of how much work it takes to nudge the spacecraft's mass in each direction, see Table 3.1, which shows the moment-of-inertia magnitudes I for rotation about each of *Voyager*'s three orthogonal axes. In this application, moment of inertia can also be called mass moment of inertia or angular mass, expressed as the integral of the radius squared times the infinitesimal increments of mass:

$$I = \int_{axis}^{edge} r^2 \delta m \tag{3.1}$$

where δm is the mass variation out along radius r from the axis of rotation to its extremity.

Table 3.1. Moments of Inertia on *Voyager*'s Body Axes (as of Saturn-Uranus cruise). Adapted from [3].

Motion	About Axis	Moment I, kg/m^2
Pitch	X	4183
Yaw	Y	588
Roll	Z	3945

In the table, note that it is rotation about the Y-axis, motion in yaw, which exhibits the smallest value of I, and therefore naturally experiences the most rapid change, clearly visible in the yaw trace in Figure 3.2.

If you stand in front of the full-scale model of the *Voyager* Spacecraft in the Von Kármán Auditorium where public lectures are held on the JPL campus, facing the model square on, the gold record of messages from Earth mounted prominently on the exterior seems to point out the center of the spacecraft's mass. The craft's optical-instrument scan platform extends off to the right. The RTG boom (Radioisotope Thermoelectric Generators supply *Voyager*'s electrical power) projects out the left side. A thin fiberglass truss-work magnetometer instrument boom is collapsed inside a shiny cylindrical canister on the *Voyager* model's left side, illustrating its state at launch in 1977. Once the spacecraft is in flight, this lightweight boom extends 13 meters up to the left. The craft's width, from the scan platform to the outboard end of the RTG boom, is 8.5 meters (see images on pp 294 and 295 in Appendix A).

If you were to walk over toward your left and approach the shielded end of the outboard RTG, you'd have a view of the spacecraft similar to the view in Figure 3.1. Lifting the RTG up would apply motion about the spacecraft's X-axis, which motion is called *pitch.* If you could twist the spacecraft's attitude by turning the outboard end of the RTG boom as though it were a helm, that motion would be rotation about the Y-axis, or *yaw.* This involves the least amount of torque, since the moment of inertia in yaw is the smallest of the craft's three degrees of freedom as evident in Table 3.1. *Roll* denotes motion about the vertical Z-axis, which goes up through the center of the high-gain antenna dish and down into the auditorium floor.

3.2 The Attitude Control System

A spacecraft's attitude has to be measured and reported, stabilized, and controlled for a number of reasons. For one, a high-gain radio antenna may need to point steadily toward Earth for communications, which is usually the case with *Voyager.* Onboard instruments have to be pointed precisely toward their targets. For some observations, an optical device such as a camera may need to track a target long enough to collect sufficient light, without letting the target's apparent drift cause the image to smear while the spacecraft speeds by. So not only the correct attitude, but also precisely controlled rates of attitude change, may be required to track a target that exhibits fast relative motion, compensating to prevent image smear. And as we have seen, attitude stability is needed for guidance: firing the rocket engine to make minor corrections to the spacecraft's flight path requires keeping the nozzle pointed in exactly the right direction during the burn.

Attitude control is one of the highly refined technologies essential to interplanetary flight. Advanced software can in certain tasks seem nearly human. While it will never pass a Turing test,[2] it will be convenient and appropriate in this chapter to treat the attitude control software and hardware somewhat anthropomorphically: the attitude control system "realizes" its situation, "knows" where to find Earth, and "takes appropriate actions." Here are the basic processes that an attitude control system undertakes:

Process inputs. The attitude control system parses real-time sensory input from specialized devices on the spacecraft including instruments that observe celestial bodies, and gyroscopes that sense vehicle rotation, as well as histories of these inputs.

Account for sloshing propellant, etc. Attitude control algorithms have to account for the effects of propellants within tanks if they slosh, affecting the spacecraft's center of mass and moments of inertia. Any flexible booms the spacecraft may have will exhibit mechanical resonances that tend to wiggle the vehicle, and these forces have to be accommodated. Also, the gyroscopic effects of any spinning masses, such as reaction wheels, must be taken into account.

Estimate dynamic situation. Given all the sensory input, and algorithms to deal with modes of sloshing and vibration and spinning mass, the system estimates the spacecraft's current state of rotation — attitude control is all about rotation around one or more axes. The state of the spacecraft can only be known within bounds of its sensory and computational capabilities while the spacecraft is rotating, so we speak of estimates rather than exact determinations.

Compare with desired situation. There is always a desired state of rotation in one or more axes that has been commanded: holding steady to fire an engine, rotating so as to track a passing target of interest, or turning to communicate with Earth. The attitude control system compares its currently estimated dynamic state to the desired state and decides what to do about any difference between them.

Apply torque as needed. Based on the difference between the commanded and the currently-estimated dynamic states, attitude control issues signals that change the spacecraft's condition: for example, *Voyager*'s attitude control system directs the propulsion system to fire quick bursts from mass-expulsion devices — rocket thrusters — to modify the craft's rotation rates and orientation. The thruster-firings evident in Figure 3.2 were keeping the antenna dish facing Earth. On a different spacecraft, the attitude control system may have the option to directly operate other devices such as reaction wheels to accomplish similar tasks. We'll examine these devices shortly.

Do routine housekeeping. As do all of a spacecraft's systems, an attitude control system formulates engineering telemetry messages and passes them to the telecommunications system for relay to Earth. We saw evidence of this at the beginning of the chapter. And like other systems it accepts, parses, and executes commands that the telecommunications system receives and relays to it.

Work reliably. All attitude control system processes must function as reliably as fine clockwork. The system must run without software bugs. It must be able to monitor a host of parameters regarding its own operations, and recognize any of a number of commonly expected problems. It must be able to take corrective actions when appropriate, including switching to redundant hardware or calling for assistance from other on-board systems. When problems do occur, the system must be able to collect all pertinent information about the problem and be prepared to issue a report to controllers on Earth. The system must be able to request the spacecraft's central computer to configure the vehicle to a known safe condition, and await further instructions from Earth. It must also be able to operate in a *critical mode* which would allow the spacecraft to continue executing a mission-critical task, such as an orbit insertion, at all cost.

Recognize anomalous torque. This is one of the many conditions an attitude control system must watch out for. If an attitude-control thruster valve were to stick open, perhaps due to some foreign matter preventing full closure, the system will sense the resulting constant torque, perhaps after counts of thruster firings to counteract it exceed a nominal value. It will have to recognize the problem, and take appropriate actions. This could mean directing the propulsion system to swap to its backup branch of plumbing to correct the problem.

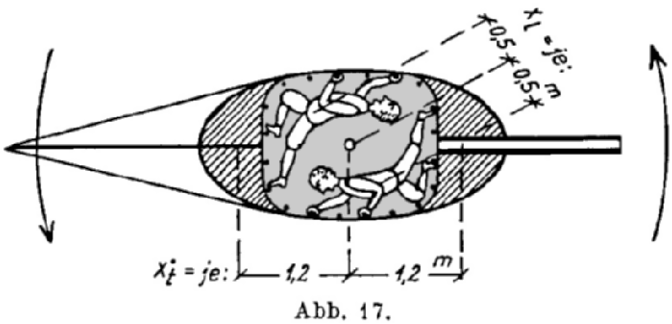

Fig. 3.3. In 1925, the German engineer Walter Hohmann (1880–1945) realized some means of attitude control would be required for spacecraft. He envisioned a system of handholds inside the vehicle that the crew could use to rotate it. Image adapted from [4].

Meet other demands. The attitude control system is called upon to serve many demands. It is often expected to satisfy a large fraction of the requirements that a spacecraft's overall design has to meet, and with some spacecraft, this can require exceptional technical achievements. For example, of all the extraordinary technical challenges that faced the *Hubble Space Telescope*'s design and implementation, including its optics, meeting the requirements on its *pointing control system* was widely viewed as the most difficult. The following are among the demands an attitude control system may have to respond to:

1. Control the motion of various articulated appendages on the vehicle, such as scan platforms, which point optical instruments toward their targets, or gimbal actuators that adjust the craft's main rocket engine nozzle pointing, or solar arrays that track the Sun and keep the spacecraft's batteries charged. It is for this reason that the system is often known as an Attitude *and Articulation* Control System (AACS).
2. Know where the Sun is. For a spacecraft whose panels must track the Sun and keep an electrical current flowing to sustain the vehicle's operation, this is a crucial task.
3. Maintain thermal control. Knowing where the Sun is in relation the spacecraft's state of rotation enables it to manage where sunlight and shadow fall on the vehicle, and keep the thermal state of the spacecraft and its instruments within predetermined limits. As an example, during their inner solar-system cruise phases, both *Galileo* and *Cassini* had to be constantly protected by keeping built-in sunshades (*Cassini*'s HGA served as one) facing the Sun. The Mercury-orbiter *Messenger*'s ability to function depends directly on such shading. And some science instruments on the spacecraft may have radiators that cool their optical detectors. They do so by facing deep cold intergalactic space, and they may not be able to withstand much direct sunlight.
4. Avoid burning the optical detectors. AACS has to comply with rules programmed onboard, such as never to point an instrument aperture within a certain number of degrees of the Sun, lest its optics concentrate sunlight onto sensitive detector electronics and fry them.
5. Know where the Earth is. Normal communication requires a high-gain radio antenna be squarely aimed toward home, and if this ability is lost only low-rate rudimentary communication is possible.[3]
6. Know where all the targets of interest are. An advanced AACS can keep tabs on the locations of any number of celestial bodies including a planet of interest, its natural satellites, as well as the Sun and Earth. This lets human controllers use a relatively high-level of commanding, such as the equivalent of "point the cameras to the center of Iapetus," instead of having to spell out precise targeting coordinates by hand, as less-capable systems may require. To implement this, AACS maintains knowledge of the bodies' motions and computes their positions out through time, using a built-in software engine called an *inertial vector propagator.*

Realtime and later: AACS's tasks are important in real time, when the craft must keep itself in the correct attitude, pointing its instruments accurately as targets come and go. In addition, a history of all the spacecraft's attitude changes supplied by the AACS serves an important function in later ground-based reconstruction of instrument pointing and spacecraft trajectory, as scientist teams proceed to analyze the results of their observations. A history of thruster firings made under AACS control is important telemetry for use in navigation, as we saw in the previous chapter.

It may be of interest to note that *Voyager*'s AACS reprogrammable flight computer accomplishes all its tasks using 4K words of memory — 8K counting the prime and usually inactive backup computer.

In the following sections we'll have a look at many of the ways various spacecraft employ AACS, we'll examine the system's many linked devices and disciplines, and we'll touch upon the propulsive capabilities, which are discussed at greater length in the next chapter.

3.3 Intersecting Disciplines

Expertise in the field of spacecraft dynamic attitude control spans several disciplines including control theory, rocket propulsion, orbital mechanics, and astronomy, as well as the enabling mathematics and physics that are ubiquitous throughout space flight.

Control theory: Spacecraft attitude control is of course primarily a control system. A simple example of control theory can be found in an automobile's cruise-control system, whose task is to keep an eye on the speedometer and issue adjustments to an *actuator* that moves the engine's throttle. Given a desired *reference* speed target by the human, the cruise control sends *output* signals to the automobile's throttle, while obtaining *feedback* information from the speedometer about the system's condition. It varies its control output until the difference between reference and speedometer, called the *error signal*, is minimized. Figure 3.4 illustrates at a high level the basic closed-loop system that applies to automobiles and spacecraft. Cruise control and AACS each utilize the *closed-loop* architecture illustrated there. Inputs from body states affect system outputs. The results of those outputs are monitored, generating an error signal that feeds back into the control algorithm.

As a basis for comparison, an *open-loop* control system is much less sophisticated. For example, a cruise-control system of decades past consisted merely of a direct mechanical friction-locked throttle position-holding knob. Start driving up or down a hill, and the open-loop system fails to maintain control of the vehicle's dynamic state. The human observes the error and then has to provide the control-system feedback by readjusting the lever.

The Scottish physicist James Clerk Maxwell (1831–1879), who is widely known for his contributions to our understanding of electromagnetism, conducted what is perhaps the first formal analysis of a control system in 1868 [5]. His study of the dynamics of a mechanical engine-speed governor helped him see how to remedy the phenomenon of "hunting," wherein he traced surges and unstable behavior to the lags inherent in the mechanical feedback. The Wright Brothers succeeded in their achievements in controlled gliding flight in 1900, and powered flight in 1903, largely because they had correctly reasoned that any free-flying object would need a control system to manage the craft's roll, pitch, and yaw. For their machine, they

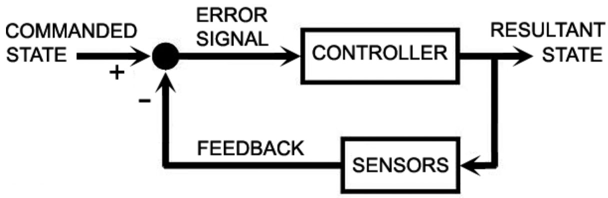

Fig. 3.4. The basic functions of a closed-loop control system. Arrows indicate data flow. Negative feedback from sensors combines with input representing the desired or commanded state to control the dynamic behavior of the system. An example might be taken from an automobile's cruise control: given a commanded state of 100 kilometers per hour, and data from the sensor showing the vehicle's speed to be 60 kilometers per hour, the error signal would tell the controller that an additional 40 kilometers per hour is required.

developed a system of moving the surface of the airfoils ("wing warping"), which the pilot could use to control the craft's attitude. Later, development of control theory [6] became important in World War II for weapons-fire control, leading to further evolution of guided missile controls and eventually space flight.

Rocket propulsion: For a spacecraft in free-fall, control theory can interface with rocket science when an AACS needs to apply a torque[4] to change the spacecraft's rotation rate. The way in which AACS interfaces with thrusters in its output path is analogous to the way a cruise-control system interfaces with the automobile engine's throttle. Components of the propulsion system accelerate and expel mass from an onboard supply in controlled directions and amounts, applying Newton's third law to answer AACS's call for torque.

Orbital mechanics: In turn, rocket propulsion under AACS's control interfaces in a couple of ways with orbit or trajectory control and determination, aspects of the art of navigation that we surveyed in the previous chapter.

First, the use of thrusters for attitude control is usually designed to produce a balanced force when applying torque to a spacecraft. For example, applying a roll torque to a spacecraft may mean firing two thrusters, each one on an opposite side of the craft, expelling mass in opposite directions. If only one thruster were to fire, the spacecraft would still feel a rotational torque about its roll axis, but the unbalanced thruster's force would also translate into nudging the whole vehicle somewhat, affecting its trajectory. Slight imbalances always exist in propulsion systems due to differences in thruster efficiency, impingement of a plume on part of the spacecraft, or nozzle misalignment, so attitude control using thrusters must always be accounted for in the navigation process.

Second, when intentional course corrections are carried out, AACS is centrally involved in directing the thrust vector in the proper direction, and managing the vehicle's attitude throughout the burn period. The AACS on some spacecraft also uses an accelerometer to determine when to cut off thrusting. As we saw at the end of the previous chapter, the only time in which a spacecraft's attitude relates to its path through space is when propulsion is used.

Astronomy: A spacecraft's intrinsic body axes of pitch, yaw, and roll, must be reckoned with an external reference frame in order to be able to estimate and control the spacecraft's interactions with the outside universe.

The first way the field of astronomy intersects with attitude control is in providing the external reference frame. Attitude control systems commonly use the reference frame defined by the standard epoch J2000.0 mentioned in the previous chapter. The spacecraft's attitude, then, is described by expressing the relationship between its own internal reference frame and the equator and equinox of J2000.0.

The relation of the spacecraft's orientation to the external, astronomical reference frame can be represented using a variety of methods that handle three-dimensional rotations. Figure 3.5 illustrates as one example the three Euler (pronounced "oiler") angles, named for the prolific Swiss mathematician Leonhard Euler (1707–1783). This and additional methods, including quaternions, are discussed in reference [7].

The second way astronomy intersects with attitude control is in the workings of appliances such as Sun sensors, which measure the apparent position of the Sun, and various devices that reckon star positions, called star trackers, star scanners, and stellar reference units (the branch of astronomy that deals with precise positions and motions of stars is *astrometry.*) All these *celestial reference* devices, each of which we'll examine later in the chapter, provide inputs to AACS for it to use in estimating the spacecraft attitude in relation to the external reference frame. Some of the latter devices achieve recognition of the "fixed" distant stars by color and brightness, or by reckoning their patterns in the sky.[5] Modern stellar reference units may contain built-in catalogs of thousands of stars including their positions, brightness, colors, and variabilities.

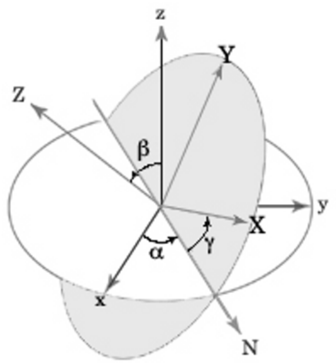

Fig. 3.5. Euler angles α, β, and γ express the relative orientation of two coordinate systems, one fixed, labeled xyz, and one rotated, XYZ. The line of intersecting nodes is labeled N.

Finally, astronomy has accumulated knowledge of the movements of target bodies of interest to a spacecraft's science investigations. Ephemerides of these natural bodies are known as a result of decades, and even centuries, of observation. And there is feedback when investigations of a target body from a precisely navigated spacecraft help refine knowledge of the body's orbit, rotation rate, and polar motions. This can be useful academically in the long term, as well as practically in the short term when optical-navigation imaging is employed to reduce uncertainty in a target's ephemeris to help negotiate an upcoming close encounter.

3.4 Stability

There are two common ways to keep your spacecraft's attitude stable. Setting the whole vehicle spinning about its central axis is one way, wherein the gyroscopic action — the inherent rigidity in space of the spin axis — of the rotating spacecraft mass about its center is itself the stabilizing force. This is a passive, open-loop means of stabilization. The other way is by using active three-axis control in a closed-loop system, as we'll see in the next subsection. Then there is the uncommon third means of maintaining stability in two out of a spacecraft's three axes while orbiting a planet, that of gravity gradient. This takes advantage of the fact that a body's axis of minimum inertial moment will naturally rotate to point toward the planet. Since the force of gravitation decreases with the square of distance, the spacecraft feels a slightly greater tug on its parts that are closer to the planet. If the orbiting body's mass is not distributed spherically, it will eventually rotate to align its axis of greatest to least mass toward the planet. The Earth's Moon, and many other natural moons in our solar system have in the same way become "tidally" locked over time, to present the same face toward the planet. This passive technique was tested in low and geosynchronous Earth orbits in the 1960s. Large enough attitude oscillations persist, so that this technique cannot meet the requirements of most modern spacecraft. Some student-developed Earth orbiters do use the technique, though, by extending a boom six meters or so in length with a small mass at its end which ends up pointing toward Earth.

3.4.1 Going for a Spin

Examples of spacecraft using the simple spin-stabilization method are numerous, and they include the *Voyager*'s predecessors *Pioneer 10* and *Pioneer 11* whose missions in the 1970s were to venture beyond Mars for the first time, through the main asteroid belt, and past Jupiter and, for *Pioneer 11*, Saturn. For such an ambitious foray into the deep outer solar system, it made sense to keep things as simple as possible, and spinning the spacecraft for stability was the best choice. A spinning platform, though, is not ideal for operating a camera that must be pointed steadily at one spot, so the Pioneers' optical instruments were designed to look radially outward and build up images line by line, scanning a narrow slice of the whole local sky in a circle as the craft flew, spinning, by its targets.

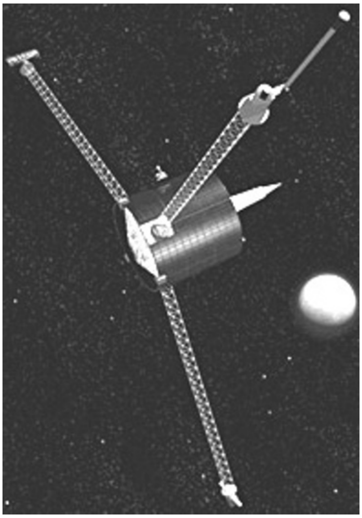

Fig. 3.6. The spin-stabilized *Lunar Prospector* spacecraft. Image courtesy NASA/Ames.

Scientists who measure the particles and the magnetic and electric fields surrounding planets, and the fields and particles in interplanetary space (and probably interstellar space too, thanks to *Voyager*), and those who wish to sample other aspects of a spacecraft's immediate environment, usually prefer to have their instruments constantly sweeping the local medium. So for them, a spin-stabilized craft is a fine platform.[6] It was a natural choice for *Lunar Prospector*, a mission flown by the NASA Ames Research Center in Mountain View, California. Reference [8] tells its whole story. This spacecraft collected data from the Moon using no cameras or other optical science instruments. Figure 3.6 depicts this spacecraft, whose radial symmetry is obviously designed for spinning. *Lunar Prospector* carried out much of its sensing of the lunar environment and surface employing a total of five instruments, mounted at the ends of three radial booms, that by design were most effective when being swept around constantly. The spacecraft rotated at 12 rpm about its Z-axis.

Measurement from Earth of the spacecraft's fine-scale changes in velocity, revealed by the Doppler shifts in its two-way phase-coherent radio signal, helped *Lunar Prospector* map the lunar gravity field and thereby characterize the distribution of surface and subsurface lunar mass. Gravity field mapping is an objective well suited to a spin-stabilized craft. By comparison, a three-axis stabilized craft's velocity is often affected by thruster firings, masking the accelerations induced solely by the gravity field under study. Spinners need propulsion systems and rocket thrusters too, to set spin rate, and perhaps to change the direction of the spin axis. But their thrusters are typically commanded to operate deliberately, instead of automatically, and only once in a long while.

Pointing cameras and other devices from a spinning platform presents challenges. The first spacecraft to orbit Venus, *Pioneer 12* (also known as *Pioneer-Venus 1*),[7] was launched in 1978 and returned data from Venus orbit until 1992. This cylindrical spacecraft carrying seventeen scientific experiments was spin-stabilized, but the great distance from Venus to Earth required it to use a one-meter diameter high-gain antenna to maintain communications. The spacecraft's design met this demand by mounting the HGA above the body center along its Z-axis, and constantly rotating it opposite the spacecraft's spin (approximately 15 rpm) using an electric motor, keeping it "de-spun" and trained on Earth throughout its flight.

The European Space Agency's *Ulysses* spacecraft, launched in 1986, operated well into 2009 in a unique high-inclination orbit about the Sun, 80° to the ecliptic plane (it attained this inclination using a Jupiter gravity-assist flyby), on a mission to characterize the heliosphere as a function of solar latitude. This highly successful spin-stabilized spacecraft had no cameras or other optical instruments, but it made many fundamental discoveries. One science experiment it carried, though, turned out to be a bit troublesome, because the spin affected a 7.5 meter-long boom. This component of the radio and plasma wave science instrument extends directly out along the spin axis, on the side of the spacecraft opposite the HGA. Uneven solar heating at certain portions of its solar orbit, combined with the boom's non-rigid mounting system, caused the axial boom to flex and impart an unacceptable amount of nutation to the spacecraft — a dynamic instability, which if left unchecked, would cause the axially-mounted HGA to wobble off Earth-point and lose contact as the spacecraft continued to spin and nutate. Specially developed procedures, involving periods of continuous uplink for over a dozen weeks at a time from the busy DSN and other facilities, succeeded in keeping nutation under control. This special procedure required programming the spacecraft to "watch" the Earth's relative position as a function of received uplink signal strength and spin rotation angle. *Ulysses*'s Attitude and Orbit Control Electronics system then fired a thruster once every three rotations to actively counteract and damp out the nutation. If the uplink were to be interrupted at the wrong time, though, the nutation could have resulted in loss of the mission. This active control of a spin-stabilized spacecraft represents an unusual case, but it attests to the ingenious capabilities that can be programmed into an attitude control system in flight.

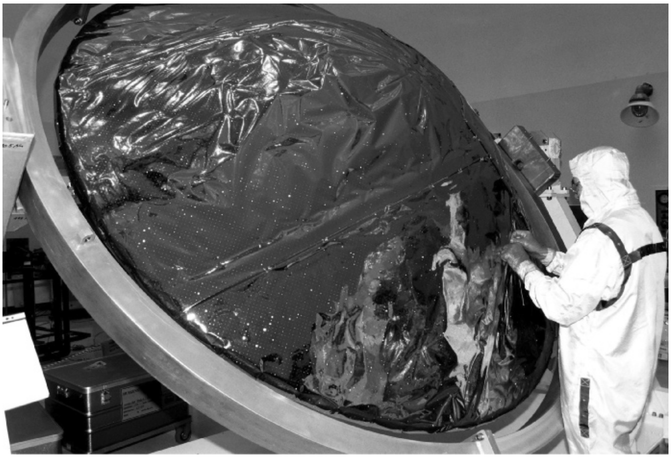

Fig. 3.7. The European Space Agency's *Huygens* Probe is a spin-stabilized craft. Here it is receiving an application of multi-layer insulation in the Kennedy Space Center's Payload Hazardous Servicing Facility six months before launch. The probe's 7-rpm spin was imparted during release from *Cassini* in December 2004. Image courtesy NASA/KSC.

Five months after arrival in Saturn orbit in 2004, the *Cassini* spacecraft was placed on a trajectory that would have it impact Saturn's moon Titan. A sequence of commands executing aboard the three-axis-stabilized spacecraft requested AACS to rotate it to a specific, pre-planned attitude. Once the attitude was achieved, *Cassini* then ejected the European Space Agency-built *Huygens* Probe it had carried from Earth, even though it was very distant from Titan at the time, and still climbing out to apoapsis in Saturn orbit. Upon release, three compressed 300 N springs expanded to push it away at 0.3 meters per second. As it departed, a curved track and roller system started *Huygens* slowly spinning, ensuring that its pre-planned attitude would remain unchanged. *Cassini* recovered from the reaction torque, turned to photograph its 319 kilograms projectile, then later it executed an OTM to avoid colliding with Titan along with *Huygens*. Before release, the

Huygens Probe had been aligned to the precise attitude that it would need to properly engage Titan's atmosphere with its heat shield without burning up. *Huygens* maintained this precise attitude, due only to its 7-rpm spin, for three weeks as it continued to orbit Saturn. When it finally slammed into Titan's atmosphere it executed a flawless descent on parachutes through the mysterious haze.

Similarly, the NASA *Galileo* spacecraft imparted a spin of 10.5 rpm to the atmospheric probe it carried to Jupiter, before releasing it on July 13, 1995. Its spin-stabilization preserved the probe's optimum angle of attack for nearly five months of free-fall until its successful atmospheric entry and descent on December 7 of that year.

Many interplanetary craft undergo a period of spin-stabilization during their launch phase. Typically, a three-axis-controlled launch vehicle places its payload in low-Earth orbit. Then, just before a powerful upper-stage rocket ignites to inject the spacecraft on its interplanetary flight, dedicated rocket thrusters fire to spin up the combined spacecraft and upper stage like a fireworks pinwheel. This provides stability while the injection burn proceeds. The 1,420-kilogram *Dawn* spacecraft, launched in 2007 to destinations in the main asteroid belt, was spun up to 46 rpm along with its attached 2,220-kilogram third-stage solid-rocket motor, to ensure attitude stability for the duration of the 87-second burn. Following this, the spacecraft needed to reduce its spin rate to near zero so that the craft's three-axis stabilization (similar to *Cassini*'s) could take over for the duration of its flight. For this it was equipped with a pair of so-called yo-yos, a commonly used system.[8] Once released, centrifugal force from the rapidly spinning spacecraft flung the two 1.4-kilogram metal masses radially outwards on 12 meter long cables that had been wrapped around the vehicle. They were let go at the end of their travel. In the four seconds this procedure lasted, the vehicle's angular momentum was literally dumped overboard, de-spinning the spacecraft and its spent, soon-to-be-detached, solid rocket. This same principle is at work when an ice skater extends his arms to stop twirling.

3.4.2 Three-axis control

As the alternative to spin stabilization, a spacecraft may be designed for active three-axis stabilization, which is the category of system *Voyager* uses. This approach is more complex and more expensive than spin-stabilization, but it offers a more maneuverable platform for pointing sophisticated optical instruments, aiming communications antennas, carrying out TCMs and OTMs, and for undertaking special operations such as described for *Cassini*'s release of the Titan probe.

At a high level, the capabilities needed for a spacecraft's basic three-axis stabilization system are:

1. A way of continuously sensing and estimating the angle between each of the spacecraft's three body axes and the external reference frame, and its rate of change;
2. The ability to determine the difference between the commanded state of rotation about each of the three axes and the observed state;

3. Some means of applying torque to the spacecraft that can rotate it in positive and negative directions about each of its three axes.

While all of these components are applicable to three-axis stabilized craft, some of them may also apply to spin-stabilized vehicles discussed above. The most important difference is that continuous automated attitude control activities are largely relinquished in the typical spinner, in favor of enjoying the built-in gyroscopic stability of the spinning spacecraft mass. Having noted this, we'll proceed to discuss the system of three-axis stabilization in particular.

Referring to the simplified closed-loop control system model depicted in Figure 3.4, we can interpret capability No. 1 in the list above as being indicated by the "sensors" box in the figure. Capability No. 2 above points to the figure's black circle combining feedback with the commanded state. Capability No. 3, applying torque to the spacecraft, would be seen as the output of the controller in the figure accomplishing the "resultant state." Within the "controller" box, and in the combiner circle, sophisticated algorithms run to compute estimates of the spacecraft's three-axis states of rotation, compare them with the external three-dimensional reference frame, generate the output signals dictating control torques that need to be applied, and watch out for potentially problematic or catastrophic situations — all the while producing telemetry and responding to command.

3.4.3 Hybrids

In summary, there are advantages and disadvantages to both spin stabilization and three-axis stabilization. Spin-stabilized craft do well with fields and particles instruments, but they may need complicated electro-mechanical systems to de-spin antennas or optics that need to point steadily at one spot. Problems with nutation can also arise. Three-axis stabilized craft can point optical instruments and antennas with ease, but they may have to carry out special rotating maneuvers to best utilize their fields and particle instruments. If thrusters provide the stabilization, observations must be designed knowing that the spacecraft is always rocking back and forth, perhaps unpredictably (to wit *Voyager*'s constant motion in Figure 3.2).

The *Galileo* Jupiter-orbiter spacecraft, launched after many delays on October 18, 1989, was designed to spin continuously for attitude stabilization. Mechanical devices on each of its three radial equipment booms could be adjusted to minimize nutation by varying the boom's angle forward or aft slightly along the Z-axis. *Galileo*'s cameras, other optical instruments, and a radio antenna for receiving signals from its Jupiter atmospheric probe, had to be precisely pointed. These requirements drove implementation of a dual-spin capability that turned out to be very complex. The lower half of *Galileo* hosting the optical devices was rotated by electric motors in the anti-spin direction, at precisely the 3 rpm nominal spin rate, to permit stable pointing. This arrangement meant devising a means for transferring electrical power and data communications across the constantly moving spin bearing. While generally successful, the feat was sometimes troublesome during operations. For some periods, an all-spin mode was needed, for example prior to probe release, in which the de-spin motor was commanded to stop. When this was done, *Galileo*'s computers experienced repeated resets, a problem that was traced to momentary interruptions in the power and data commutators when there was no relative motion across the bearing. The remedy was to create a "quasi-all-spin" mode that kept the de-spun section moving very slowly to help the commutator maintain electrical contact without interruption.

The *New Horizons* spacecraft (see page 296 in Appendix A) is using spin stabilization for much of its cruise out to Pluto and other Kuiper Belt objects. During launch from Earth early in 2006, its spin rate was increased to 68 rpm for maximum stability while its solid-fuel rocket motor burned, with characteristic unevenness, to inject the vehicle onto its fast interplanetary trajectory. Then after injection, its spin was reduced by releasing yo-yo weights to 5 rpm for the long haul past Jupiter and on to its intended targets. During its planned encounters, *New Horizons* will stop spinning and go into three-axis stabilization mode, as it also does during periodic checkouts en route.

3.5 Attitude Control Peripherals

There are a number of items under the category of input devices, the *sensors* that gather information about the state of the system being controlled. And there are the various *actuators,* the output devices that an AACS uses to exercise its control over the system. Broadly, AACS sensory inputs come from either celestial or internal reference devices. Its use of output devices applies torque to the spacecraft in various ways, bringing its attitude and rotation rates into conformity with commanded states.

3.5.1 AACS Input Devices

Celestial Reference

A *Sun sensor* is a common AACS celestial-reference — sky-watching — input device. It is an optical sensor with a wide field of view that reports on movements of an image of the Sun in two axes across its light-sensitive detector. The traces of *Voyager*'s excursions in yaw and pitch in Figure 3.2 on page 88 are readouts from a Sun sensor. Typically, spacecraft have at least two of these important devices for redundancy in case one were to fail. For a *Voyager* or *Cassini*, whose Sun sensors have a view along the Z-axis, the devices are sensitive to spacecraft attitude changes in two degrees of freedom, pitch and yaw, and they report these to ACCS. They do not sense activity in roll.

The large parabolic reflector of *Voyager*'s High-Gain Antenna, HGA, is usually facing back toward the Earth, which is nearby the Sun as seen in *Voyager*'s sky in the far reaches of the solar system. The HGA was therefore designed with a hole in it, through which the Sun sensors have a view toward the inner solar system.[9] In April 2002, engineers switched off *Voyager 1*'s primary Sun sensor, and activated the backup. After twenty-five years in flight, it had begun showing some signs of degradation.[10] On *Cassini*, Sun sensors occupy two holes through the spacecraft's HGA, widely spaced so that attitude control could be maintained in case a stray ring particle were to damage either the prime or the backup Sun sensor while the spacecraft orbits Saturn. The European Space Agency's *Mars Express* spacecraft, orbiting Mars since late 2003, has two Sun sensors, one of which was used for initial attitude determination following launch.

Another kind of device on the typical interplanetary spacecraft looks off approximately at right angles to the Sun sensor's view, to provide additional reference information by observing one or more background stars. *Star-watching devices,* as with many components on a spacecraft, are usually present in a redundant pair providing for backup in case one were to fail. On *Voyager*, the Canopus Star Tracker, named for the single bright star it was designed to watch, provides measurements in the one remaining degree of attitude freedom: excursions in roll. Measurements from this device are seen in the bottom panel of Figure 3.2 (page 88) as they are reported to AACS. *Voyager*'s Canopus tracker can be trained on other bright stars besides Canopus, by rolling the spacecraft, although only one star can be tracked at a time.

Fig. 3.8. The backs of *Voyager*'s sun sensors are visible on the white HGA above the heads of the people affixing the famous golden record. Two of the four yaw thrusters can be seen below the record. Image courtesy NASA/JPL.

Somewhat more advanced than *Voyager*'s single-star tracking device, a "V-slit" star scanner provides complete attitude reference while affording more freedom of motion. Three-axis-stabilized craft that use these devices must execute a rotating maneuver to obtain a star-scan attitude reference, while spin-stabilized craft can use them for continuous reference. See Figure 3.10 and we'll explore how it works. The scanner views the background of stars through two slits that are not parallel to one another. As the spacecraft rotates, the appearance of a star in the first, vertical slit, produces a voltage proportional to the star's intensity, called a "clock" signal. The time at which the same star passes through the next slit, the slanted one, marks the "cone" signal. After accumulating a number of these events in memory, the tracker's built-in computer algorithms, referring to an internal database of star position and brightness information, can proceed to deduce the spacecraft's attitude. The spinning *Galileo* Jupiter-orbiter spacecraft used this kind of device, as did the three-axis-stabilized *Magellan*.

More mature in design than the single star-tracker or the V-slit scanner is an autonomous Stellar Reference Unit (SRU). Two of these devices are fixed to the *Cassini* spacecraft's side, looking orthogonally to the Sun sensor's view (see page 330 in Appendix B). The SRU is not constrained to view only one star, nor is it constrained to view a moving star field. It observes the entire field of stars in whatever direction it is pointing, and it accomplishes recognition of a number of them based on the stars' observed geometry and intensity, by comparing them against its built-in catalog. It can do this whether the spacecraft's attitude is changing or not. This sophisticated device provides attitude reference for all three axes at once. A high-performance modern SRU may have a square field of view 8° or so on a side, be able to recognize and track a dozen stars at once, with onboard knowledge of thousands. It can be expected to report the observed spacecraft attitude to AACS in reference to the J2000.0 inertial reference frame with high accuracy.

Star-watching devices are sensitive instruments. They can be confused if a bright nearby object such as a planet or a ring system enters their field of view, or if its view is blocked by the night side of a planet. Mission commanding sequences must therefore tell the AACS in advance to ignore input from such devices for periods when such an obstruction may be present, or else attitude knowledge may become corrupted. For *Voyager*, simply executing a turn may result in an attitude from which all celestial reference — the Sun and one background star — is lost.

Usually, before a star-watching device can begin to recognize stars and provide reference information to AACS, the Sun must be visible within the Sun sensor, narrowing down possibilities for the spacecraft's attitude and providing an important initial scenario. V-slit scanners and SRUs can then continue to provide reference data after the Sun has left the Sun sensor's view. In anomalous situations when a spacecraft's AACS has lost all attitude

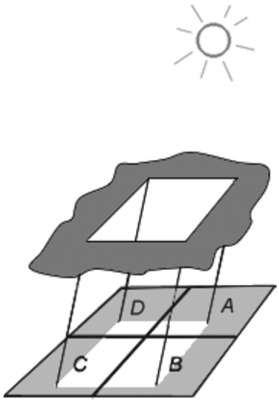

Fig. 3.9. In one type of sun sensor, four rectangular photovoltaic cells, A, B, C, and D, receive varying amounts of illumination based on incident sunlight falling through a rectangular aperture centered above them. If the sun-line were normal to the sensor, all four cells would have the same amount of illumination and would output the same electrical signal.

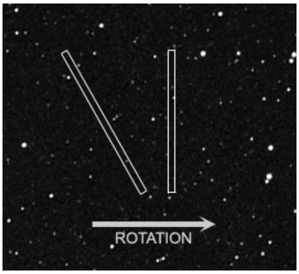

Fig. 3.10. Artist's conception of a V-slit star scanner's view against a field of stars. Spacecraft rotation first passes a star through the vertical "clock" slit as the slits move to the right. The star's subsequent passage through the slanted "cone" slit, based on a-priori knowledge of star-field geometry, provides enough information to determine spacecraft attitude in three axes.

knowledge, a typical autonomous protective response is to execute a maneuver that rotates the spacecraft, sweeping the Sun sensor's field of view around the 4π-steradian sky until the Sun is re-acquired.

Some orbiter spacecraft also carry a *Horizon sensor*. This optical instrument detects visible or infrared light from the planet's limb, or from its atmosphere, and provides information on the spacecraft's orientation with respect to the planet about two orthogonal axes.

Inertial Reference

Self-contained attitude reference devices that do not depend on external input are needed since celestial reference devices such as Sun sensors, star trackers, star scanners, stellar reference units, and horizon sensors, cannot be used under all conditions during a mission, as noted above. For such times, for example when a spacecraft passes into the shadow of a planet, a spacecraft's attitude control computer may need to have an independent reference. *Inertial reference* inputs are generated by *angular-velocity sensing* devices, known as gyroscopes (from the Greek words meaning "rotation" and "to see"),[11] or "gyros" for short, which do not depend upon making any observations outside of themselves.

There are a number of mechanical principles that can serve as the basis for gyroscopes. A small, rapidly spinning mass can be readily used because of its gyroscopic property of rigidity in space stemming from the mass's angular momentum. Employing a set of low-friction gimbals, a spinning-mass-gyro-based Inertial Reference Unit (IRU), is able to measure the apparent rotation of the gyro, which is largely fixed in space, as the spacecraft basically rotates about it. *Voyager*'s IRUs employ spinning-mass gyros, whose output provides rates of spacecraft rotation. Typically, a complete inertial reference unit uses three gyros, one each to sense excursions in pitch, roll, and yaw. These devices have been widely used in aviation for decades.[12]

Don't throw away your celestial reference devices yet, though. Spinning-mass gyroscopes are imperfect attitude references, because there is always some frictional coupling between their motor-driven internal spinning masses and their gimbaled mounts within an IRU. So they precess. The result is that the reference signals they produce typically drift, and exhibit errors that build up over time in reckoning the spacecraft's true attitude. Gyros that use different physical properties, which we will see below, also suffer from inaccuracies, even though they may not be subject to friction. Inertial references, then, are typically called upon for relatively short periods when celestial references cannot be used. To be useful, an IRU's errors have to be calibrated in flight using celestial references. Once an IRU's drift rates are known, they can be routinely compensated for by commanding the IRU to update stored drift-bias values regularly from Earth following calibration maneuvers. The *Hubble Space Telescope*, for example, requires this procedure to be done once every several days. Some spacecraft only use their gyros for infrequent maneuvers, so their drift calibrations may be carried out just prior to each use.

The *NEAR-Shoemaker* spacecraft, launched on its near-Earth asteroid observing mission on February 17, 1996, and the *Cassini* spacecraft orbiting Saturn, are the first interplanetary craft to use gyros that have no spinning parts. *Messenger*, launched in 2004 for Mercury, and many Earth-orbiting craft also use them. Their inertial reference elements, known as Hemispherical-Resonator Gyros (HRGs), operate on a different physical principle from the gyroscopic rigidity in space. These intriguing devices sense movement of a standing mechanical wave induced in the rim of a fused-quartz (crystalline silica, silicon dioxide SiO_2) hemispherical shell about 3 centimeters in diameter. The wave is akin to that in a crystal wine glass ringing like a bell when struck (see Figure 3.11). Null points in the wave travel about the rim at a different rate than the glass itself, when it is rotated about its axis of symmetry. The British professor George Hartley Bryan (1864–1928) first described this principle in 1890 [9]. The feat HRG devices accomplish is to induce a continuous ringing-vibration in the hemisphere, and detect and track the null points' motions with great sensitivity, by taking advantage of its piezoelectric[13] properties. Other than their vibrating sensor shells, hemispherical resonator gyros have no moving parts, and have nothing to wear out.

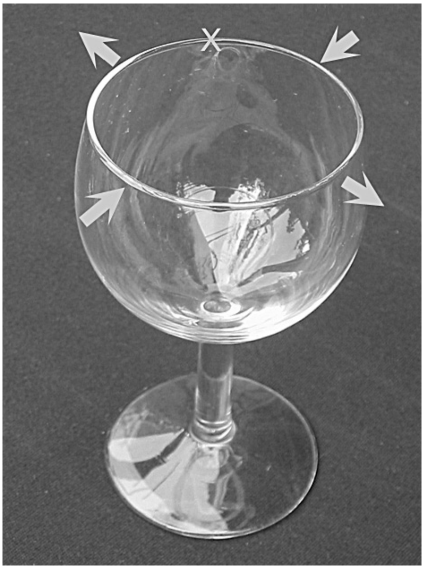

Fig. 3.11. A wine glass serves as an analog for a hemispherical resonator gyro. If the glass is made to ring audibly, a snapshot of rim dynamics would show flexing as indicated by arrows, which periodically reverses. Nodes between arrows such as X exhibit minimum flexing. These nodes precess about the rim at a different rate than the glass itself when rotated about its vertical axis.

Laser gyros are commonly used in aviation applications and are employed on some spacecraft. They use the Doppler shift of light to sense attitude rate changes about each axis in which they are mounted. The *Clementine* spacecraft, which orbited the Moon in 1994, employed these devices, as does the *Mars Express* spacecraft. Two light beams are sent in opposite directions along a medium in one plane — either a fiber optic line, or vacuum and mirrors. When the system is rotated in-plane, light going along one path travels farther than the light going in the opposite path during transmission, as seen in the familiar Doppler effect. This causes the light waves to interfere with one another, producing measurable patterns known as Sagnac Interference, named for the French physicist Georges Sagnac (1869–1926) who studied the phenomenon and identified its cause. Spaceborne systems usually use several kilometers of optical fiber wound in a coil for each of the three axes of rotation to be measured.

Micro-Electro-Mechanical Systems (MEMS) gyros[14] use another principle. MEMS gyros, produced using the same silicon etching processes that are used to make electronic chips, employ tiny, rapidly vibrating flexible arms. The principle at work is the same that we observe in a Foucault pendulum: Vibrating or oscillating objects tend to continue moving in the same plane. Rotating the system results in a Coriolis-effect[15] torque that can be measured. MEMS gyros typically use the piezoelectric effect to keep their test masses vibrating, as well as to generate an error-signal voltage proportional to rotation. Also known as "ceramic gyros," the inexpensive devices are found in today's consumer electronics including digital cameras to provide image stabilization, hand-held 3-D computer input devices that control cursor position or game components, and the Segway® Personal Transporter. NASA's New Millennium Space Technology-6 program included the launch of *TacSat 2* into low Earth orbit in December 2006. This small spacecraft demonstrated the use of an integrated SRU and three-axis MEMS gyro set for attitude control reference, called the Inertial Stellar Compass. This compact, low-power package that combines celestial and inertial references has a mass of 3 kilograms and draws only 3.5 watts from the spacecraft electrical supply.

There's one more kind of inertial reference device spacecraft carry to send input to AACS. On *Cassini* and other spacecraft, an *accelerometer* provides measurements of the force applied to the spacecraft during rocket engine burns for TCMs and OTMs. In most cases, AACS parses accelerometer input to compute when to shut off the engine after it has provided a specified value of ΔV. Science instruments use accelerometers as well. The *Huygens* Atmospheric Structure Instrument, carried aboard the *Huygens* Probe (see page 327 in Appendix B), contained, among its other components for measuring temperature and pressure, three accelerometers that registered forces acting on the probe in all three axes as it descended through Titan's atmosphere. *Huygens*'s Surface Science Package also included accelerometers that measured the force of landing (15 g), as well as Titan's natural gravitational force on the surface (a little less than 1/7 g). When the *Mars Global Surveyor* and the *2001 Mars Odyssey* spacecraft were executing aerobraking maneuvers, dipping into Mars's upper atmosphere, on-board accelerometers generated data that were used to derive atmospheric density values. Atmospheric entry vehicles that carried the *Mars Pathfinder* (1997) and Mars Exploration Rovers *Spirit* and *Opportunity* (2004) to the planet's surface also reported forces experienced by on-board accelerometers. Accelerometers on the rovers themselves indicate which way is down during surface operations. Many Earth-based navigation systems use accelerometers to add up all the movements of a vehicle — for example, an airplane — and form a complete picture of the vehicle's path from point to point.

3.5.2 AACS Output Devices

Mass Expulsion We've alluded to rocket thrusters earlier in this chapter, as well as in the previous one. In the next chapter we'll look more closely at how they work. For the present, we'll consider their role as common AACS output devices. Systems employing thrusters for attitude control are also referred to as mass-expulsion control (MEC), or reaction-control systems (RCS), named for the reaction obtained from the action of expelling mass according to Sir Isaac Newton's third law. By selecting which of several MEC thrusters to use, AACS can apply torque to a spacecraft about any of its axes. Varying the amount of time thrusters apply torque will vary the spacecraft's attitude change rates.

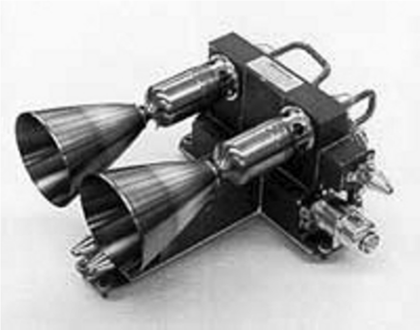

Fig. 3.12. *Magellan* rocket thrusters. The largest ones seen here developed 445 N during TCMs and Venus orbit insertion. The mid-size 22 N thrusters (right side) controlled roll while the 445 N thrusters were in use. The smallest, 0.9 N, were used for routine reaction-wheel desaturations. Courtesy NASA/JPL.

The ten *Mariner*-class spacecraft that JPL built in the 1960s were intended to explore the inner solar system. Six of them survived launch and accomplished their missions to Venus, Mars, and Mercury. These were the first interplanetary spacecraft to depart from the spin-stabilization design and use three-axis control. Their mass-expulsion devices were as simple as can be. Each *Mariner* spacecraft was equipped with a total of twelve small nozzles mounted at the ends of its four radially oriented solar panels. When the spacecraft's AACS called for a torque to be applied, it opened an electrically controlled valve for 20 ms, supplying compressed cold dry nitrogen to each of two opposing nozzles. This permitted gas to escape from a common tank, providing a thrust of about about 0.1 N from each nozzle and forcing the spacecraft to rotate. Reference [10] describes the system. In the interest of relating to familiar experience, consider the consequence of letting go of a garden hose while its nozzle is expelling water. The action of water accelerating out through the nozzle produces a reaction causing the nozzle to travel backwards.

The two *Voyagers* are *Mariner*-class spacecraft, but the mass-expulsion control system accessible to their AACS is more sophisticated than that of the previous *Mariner*s. Each *Voyager* has sixteen small liquid-propellant rocket thrusters that deliver a push of about 0.9 N each. Note that two of *Voyager*'s yaw thrusters, a prime and a backup, are visible in Figure 3.8.

Voyager's AACS operates the thrusters in pulses lasting a number of milliseconds during which an electrically controlled valve opens to spray hydrazine (N_2H_4) onto an electrically heated catalyst in the combustion chamber, which causes the propellant to decompose explosively, rapidly expelling hot gas. After encountering Uranus in 1986, the software capability to reduce each thruster pulse from 10 ms to 4 ms was developed, tested, and installed on the spacecraft in flight. This permitted finer attitude control during long camera exposures in the dimly lighted Neptunian environment (less than 1/1600 the sunlight that we enjoy on Earth), while also extending the life of *Voyager*'s propellant supply.[16]

Attitude control thrusters may be called upon to apply large torques to a spacecraft, typically while a more powerful rocket is operating to impart significant ΔV to the spacecraft. During launch, *Voyager* ignited a solid-propellant rocket motor that provided a final increment of speed to begin its free-fall cruise to Jupiter. Because solid rocket motors typically burn somewhat unevenly, they can impart strong off-center components of thrust and perturb the spacecraft's attitude. To maintain control, each *Voyager* used four 445 N monopropellant thrusters on struts straddling the solid rocket motor. Figure 3.13 shows where all the RCS engines and thrusters are mounted on the *Magellan* spacecraft in a similar arrangement to that of *Voyager*, at the ends of four struts. Because the struts reached out from below the spacecraft's center of mass, thrusters mounted there were able to overcome torques resulting from the 67-kN solid rocket's 84-second burn that resulted in an acceleration force up to 7 *g* placing *Magellan* into Venus orbit. The figure does not show the solid motor, which was jettisoned after use. Mounting attitude control thrusters out on struts increases their leverage, or control authority, since the distance out from the center of mass determines how much torque a thruster can wield on the spacecraft when it applies its given amount of force.

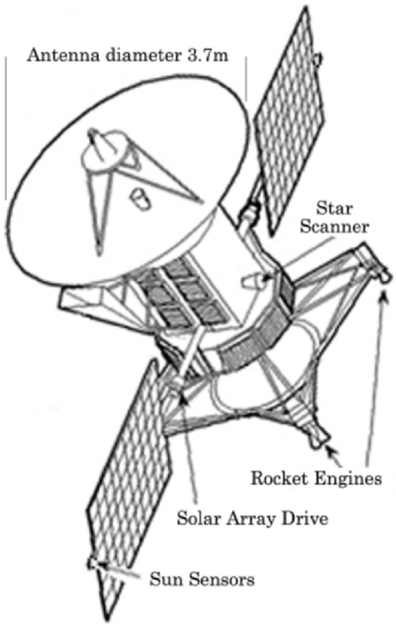

Fig. 3.13. The NASA Venus-mapping *Magellan* spacecraft. Adapted from NASA/JPL-Caltech image.

Reaction Wheels There's another kind of output device for applying torque for three-axis stabilization. Electrically powered reaction wheel assemblies (RWAs), can impart a torque under control of AACS to the whole spacecraft. Note that reaction wheels are sometimes called "momentum wheels," but the latter name is also applied to a different system, called control-moment gyros, which we will discuss separately. In the RWA system, small but fairly massive wheels are mounted aboard the spacecraft with their rotational axes fixed. *Magellan*, whose three electrically driven reaction wheels were mounted near the center of mass with their axes oriented orthogonally to one another, is a good example. To rotate the vehicle in one direction, the attitude computer causes one of the wheels to accelerate in the opposite direction and remain spinning. When the wheel has finished accelerating, the spacecraft itself has acquired a steady rotation rate. To stop the vehicle's rotation, the AACS would simply slow down the same wheel. This system provides a means to trade *angular momentum* back and forth between the whole spacecraft and its reaction wheels. In practice, RWAs that use a fluid lubricant are usually operated with some residual spin, or *bias,* to prevent lubricant stagnation near zero rpm.

Consider that a large mass, such as a whole spacecraft, changing its attitude at a relatively low rate, can have the same angular momentum as a small mass spinning rapidly. In reference to equation 3.1 on page 88, the spacecraft has a high

I while a small mass would have a low I. Angular momentum, expressed as the vector quantity $\vec{H}$, is the product of an object's moment of inertia, I, and its angular velocity, which is typically expressed as the vector omega, $\vec{\omega}$:

$$\vec{H} = I \cdot \vec{\omega} \tag{3.1}$$

Physics requires that in the absence of any externally imposed torque, the angular momentum of a whole system, such as a spacecraft containing reaction wheels, must remain constant. Adding torque to a reaction wheel, spinning it up and increasing its angular momentum, has the effect of decreasing the angular momentum of the rest of the spacecraft — the decrease can mean it begins rotating "backwards." Likewise, decreasing a reaction wheel's $\vec{H}$ will increase that of the rest of the spacecraft. The total angular momentum vector of the spacecraft at any time while under RWA control, then, will have two components (in the absence of externally applied torque). Expressed in spacecraft-body frame:

$$\vec{H}_{Total} = \vec{H}_{SC} + \vec{H}_{RWA} \tag{3.2}$$

where $\vec{H}_{SC}$ represents the component due to spacecraft angular rates, and $\vec{H}_{RWA}$ is that due to the reaction wheels. On the *Cassini* spacecraft, each electrically driven reaction wheel has a mass of 14.5 kilograms, a diameter of 30 centimeters, and a maximum speed around 3,000 rpm. These are effective in rotating the approximately 5,700 kilogram spacecraft at rates up to about 1.5 mrad/s in pitch and yaw, and twice that in roll, the axis with the smallest moment of inertia. Reference [11] gives the context, details, and performance of *Cassini*'s system. Wheels provide excellent stability and precise control for pointing optical instruments, meeting the *Cassini* science requirements ranging from 8 μrad precision for a one-second observation to 160 μrad for 100 seconds. Minutely varying the speed of a rapidly spinning small-I device affords a precise level of control on the larger-I spacecraft not unlike the way a reducing gear train offers fine-scale vernier-control of an output shaft's angle.

A good "seat-of-the-pants" way to visualize the basic mechanics of reaction wheels at work on a spacecraft is to carry out a thought experiment. Imagine holding a battery-powered electric drill while sitting, feet up, in a swivel chair. There is a circular 10-kilogram concrete paving stone with a spindle installed through its center and inserted into the drill's chuck. Keeping the concrete disk's spindle aligned vertically, you apply torque and begin to spin the heavy wheel via the drill. The result is that you and your seat begin rotating, as the reaction to adding angular momentum to the heavy wheel.[17] Now, reverse the drill-powered wheel, and you and your swivel chair will rotate in the opposite direction. One can easily imagine how vivid the results would be were one to be free-falling in orbit instead of sitting in a chair on *terra firma.*

Incidentally, the spinning masses of reaction wheels on a spacecraft do exhibit gyroscopic effects, but these are side effects that the attitude computer is tasked with calculating and working around during normal operations. Reaction wheels should not be confused with a spacecraft's mechanical spinning-mass gyroscopes, which as we have seen are *input* devices that provide inertial attitude references.

Reference gyros employ much smaller spinning masses, and their spin axes are not rigidly affixed to the spacecraft. Reaction wheels are strictly output devices that AACS uses for directly controlling attitude. Neither should reaction wheels be confused with control moment gyros, which are discussed below.

In practice, there is almost always some measure of external torque being applied to a spacecraft from solar photon pressure, gravity gradient, or atmospheric drag. These cause excess momentum to eventually build up in a reaction-wheel system as it strives to keep the spacecraft in a desired dynamic state. In its attempt to counteract a torque it senses on the spacecraft, AACS will continue to increase the reaction wheel's spin rate. Friction within a reaction-wheel system tends to cancel out excess momentum buildup, rather than contribute to it.

How can solar photon pressure affect a spacecraft's attitude? Light (and electromagnetic radiation at other wavelengths) that strikes a surface exerts a force upon it. Even though photons have no mass, because they travel at the speed of light, their energy exhibits momentum. The amount of force a spacecraft feels is related to the received energy from the Sun, which diminishes as the square of distance, and of course the amount of area illuminated. If the surface reflects light at all, it will add another component of force due to the reaction from turning it around and sending it back. The angle at which the surface faces the Sun is another factor. To estimate the total amount of solar light-pressure force:

$$F = \left(\frac{F_S}{c}\right) A_s(1 + r)\cos\Theta \tag{3.3}$$

where

F is the force in newtons,
F_S is the Sun's radiated energy in W/m^2. For example at Earth's location, F_S= 1371 W/m^2, and it is approximately 1 percent this amount at the distance of Saturn,
c = the speed of light, about 3 $\times 10^8$ m/s in vacuum,
A_s = the area of the spacecraft's illuminated surface in m^2,
r = the surface's reflectance: 0 for a perfect absorber, 1 for a perfect reflector,
Θ = the illuminated surface's angle of incidence to the Sun.

This force, although small, acts on the whole spacecraft, pushing it away from the Sun. But if there is an offset between the center of photon pressure on spacecraft's Sun-facing side and its center of mass, this will result in an external *torque* being applied to the spacecraft in a fixed direction, gently trying to rotate the spacecraft. The attitude control system senses the tiny rotation, and commands the reaction wheels to accelerate to cancel out this torque. An example is the *Mars Climate Orbiter*, that had one large solar panel to generate its electrical power, attached to only one side of the spacecraft (see Figure 2.1 in Chapter 2 on page49).

A word about the solar wind may help avoid confusion. The Sun's *light* exerts the noticeable pressure. Charged particles streaming out from the Sun, known as the *solar wind,* do not have an appreciable force on a spacecraft. Though they do have mass, they are too sparsely distributed in interplanetary space, and they travel slowly in comparison with light.

Gravity gradient can also cause a constant torque if the spacecraft is orbiting a planet. Flight through the upper reaches of the atmosphere of a planet being passed or orbited can also impose a torque, if the center of exposed area differs from the spacecraft's center of mass.

No matter the source of a constant externally generated torque, as the RWAs compensate, wheel speeds might eventually become excessive. Approaching maximum rpm is called "saturation," in which the spinning wheels are carrying as much angular momentum as their mechanical design can safely tolerate, beyond which the assembly might suffer damage.

So, to maintain wheel speeds within prescribed limits, excess momentum (excess wheel speed) must be occasionally removed from the system. This can be done by somehow applying torques to the spacecraft to hold it steady, while the attitude computer causes the wheels to slow down, typically, and acquire a desired preset speed, which may be zero, or it may be a bias of some rpm value in one direction or the other. This task is done during maneuvers variously called angular momentum desaturation (desat), reaction-wheel desaturation, momentum unload, or momentum dumping maneuvers. Many spacecraft use a system of thrusters to apply the torque needed to steady the spacecraft for desaturations. *Magellan*'s RCS thrusters were called on routinely to do this while in Venus orbit.

Magnetic Torquers The *Hubble Space Telescope*'s pointing control system uses reaction wheels to control the spacecraft's attitude. The system makes it possible to point to a target without deviating more than 0.007 arc-second — the width of a human hair viewed at a distance of more than a kilometer. Operating in Earth orbit, it is subject to relatively strong photon pressure from the Sun, plus gravity gradient from Earth, so its reaction wheels must occasionally be desaturated. But *HST*'s optics, including its 2.4-meter diameter primary mirror, are exquisitely sensitive and could easily be contaminated and rendered useless if there were rocket thrusters routinely expelling clouds of exhaust. So *Hubble* employs an alternative way to hold a steady attitude during its reaction wheel de-saturation maneuvers. The solution is to employ magnetic torquers — electromagnets in the form of four 8.5 meter-long wire-wrapped bars arrayed around the spacecraft's exterior. When energized with electric current, under control by AACS, their interaction with the Earth's natural magnetic field is powerful enough to hold the spacecraft's attitude steady while the reaction wheel speeds are modified during desaturations. Many spacecraft that operate in Earth orbit, where the magnetic field is useable (its strength at orbital altitudes is less than half a Gauss), rely on this kind of system. Tens of thousands of kilometers out, though, the field effectively ends, and torquers cannot be used.

The *Spitzer Space Telescope*[18] orbits the Sun at about the same distance as Earth does, trailing along behind the Earth in its yearly progress. As of late 2008 its distance is nearly 1×10^8 km from Earth. Reaction wheels provide steady attitude control as the telescope points toward its targets, and rotates it to point its HGA to Earth. While the spacecraft's location is convenient for making observations in deep space without the Earth getting in the way (which can often interfere with *HST*'s observations), there is no magnetic field strong enough for magnetic torquers to use during reaction wheel desaturations. *Spitzer*'s optics, designed for infrared astronomy, are even more sensitive than *Hubble*'s when it comes to contamination, because they are kept cold for infrared viewing — only 5.5 kelvins — so that its instruments can observe in the far-infrared part of the spectrum (see page 298 in Appendix A). If the spacecraft were equipped with hydrazine thrusters like *Magellan*'s, the ammonia and other products in their exhaust clouds would quickly finds ways to condense on the frigid optical surfaces and contaminate them. So to stabilize during desaturations, Spitzer issues pressurized cold dry nitrogen from nozzles, despite the relative inefficiency of such a system, a throwback to the original *Mariner* spacecraft's means of three-axis attitude control.

Spacecraft (Launch)	**Planets**	**AACS Input Devices** (dof = degree-of-freedom)	**AACS Output Devices** (Excluding any jettisoned)
Voyager-1 Voyager-2 (1977)	Jupiter Saturn Uranus Neptune (Flybys)	1-dof Canopus Star Tracker 2-dof Sun Sensors (2) 2-dof Gyroscopes (3)	0.9-N Trajectory-Correction Thrusters (4) 0.9-N Attitude Thrusters (12) Scan Platform Gimbal Actuators (4)
Magellan (1989)	Venus (Orbiter)	1-dof Star Scanner (1 with redundant channels) 2-dof Sun Sensors (2) 2-dof Gyroscopes (4)	Reaction Wheels (3) 445-N Orbit Injection Attitude Thrusters (4) 22-N Attitude Thrusters (8) 0.9-N Attitude Thrusters (12) 1-dof Solar Array Drive Mechanisms (2)
Galileo (1989)	Jupiter (Orbiter)	Sun Acquisition Sensors (4) Star Scanners (2) 2-dof Gyroscopes (2) Accelerometer (1)	10-N Attitude Thrusters 400-N Main Engine (1) Spin Bearing Actuator (1) Scan Platform Actuator (1) Linear Boom Actuators (2)
Mars Global Surveyor (1996)	Mars (Orbiter)	Mars Horizon Sensor (1) Celestial Sensor (1) Sun Sensors (2) 1-dof Gyroscopes (4) Accelerometers (4)	Reaction Wheels (4) 2-dof HGA Gimbal Actuators (2) 4-N Thrusters (12) 2-dof Solar Array Drive Mechanisms (2) 659-N Main Engine
Cassini (1997)	Saturn (Orbiter)	3-dof Star Trackers (2) 2-dof Sun Sensors (2) 1-dof Resonator Gyros (4) Accelerometer (1)	Reaction Wheels (4) 445-N Main Engines (2) 2-dof Engine Gimbal Actuators (2) 0.9-N Attitude Thrusters (16)

Fig. 3.14. Peripheral devices on the inputs and outputs of AACS for six spacecraft. Adapted from [11].

Control Moment Gyros While not applicable to most interplanetary spacecraft, we'll discuss these devices to distinguish them from RWAs. The International Space Station (ISS,) is equipped with control-moment gyros (CMGs). These are spinning-mass devices, also called gyrodynes, whose rotors are on the order of 100 kilogram mass, kept going at a constant speed by electric motors (note this difference from reaction wheels, which vary their speed). The gyroscopic properties of rigidity in space and precession are used to apply torque to the whole spacecraft. To turn the spacecraft, you rotate the *spin axis* of a CMG (recall RWA spin axes are fixed to the spacecraft body). CMGs are attached to the spacecraft structure via a set of gimbals to permit movement of their axes. Brute force precession then results in a torque applied to the whole spacecraft. The space station uses a set of four CMGs to provide controllability in three axes, keeping one as a spare in case of failure in one of the others. While CMGs have the same purpose as that of reaction wheels, note that the operating principle is different. RWAs apply torque by changing rotor spin speed; CMGs force-tilt the rotor's spin axis without necessarily changing its speed. CMGs are best suited to applications on very massive spacecraft such as today's ISS, or the *Mir* space station of the past. A set of CMGs may consume a few hundred watts of electrical power, and produce thousands of newton-meters of torque.

Another thought experiment may be appropriate to illustrate CMGs in operation. Imagine[19] sitting feet-up in your swivel chair, holding the cordless-drill-powered 10-kilogram concrete disk as in the reaction wheel thought experiment. This time, let it spin with its drill-mounted axle *parallel* to the floor. Increase its spin to maximum and keep it at that speed. Now tilt the drill, bringing its axis of rotation to an angle with the floor. Precession will cause you to rotate, just as it causes the space station to rotate.

Ancillary Actuators Attitude control is one function of AACS, articulation is the other. Following is a list describing some of the more common spacecraft components under the control of an AACS:

1. Solar array drives: Spacecraft that depend on sunlight for their electrical power supply require the photovoltaic cells on their solar panels to face the Sun. Solar array drive mechanisms have one or two axes of freedom, each operated by an electric motor. AACS maintains knowledge of the Sun's position, and can orient the photovoltaic to face it, or to employ an offset requested by the electrical power subsystem to reduce the amount of power they generate by pointing them slightly away from the Sun.
2. Engine gimbal actuators: Some spacecraft control the direction their main rocket engine's nozzle is pointing, to keep the rocket thrust directed through the spacecraft's center of mass. Based on the 1970s *Viking* Mars orbiter's design, *Cassini*'s two gimbal actuators control each main engine, constantly making small adjustments in the engine's position to compensate for shifting propellant mass, under control from AACS.
3. Scan platforms: Spacecraft that carry optical instruments on moveable platforms depend on AACS to maintain control of their pointing. *Voyager*, for example, can articulate its scan platform in two degrees of freedom. *Galileo* was able to articulate its optical instrument platform in one degree of freedom. A second degree of freedom was provided by adjusting the de-spin rate in roll of the spacecraft's lower despun section, under AACS control.
4. High-Gain Antennas: HGAs often occupy booms protruding from the spacecraft, and can be articulated in one or more degrees of freedom.

5. Linear Boom Actuators: The *Galileo* spacecraft had three booms extending radially from its spinning central body: two RTG booms and a science instrument boom. These needed to be adjusted slightly up or down along the roll axis to minimize wobble or nutation. In flight, AACS controlled linear actuators supporting the booms that were able to extend or contract up to 5 centimeters to make the necessary adjustments. These are described in reference [12].

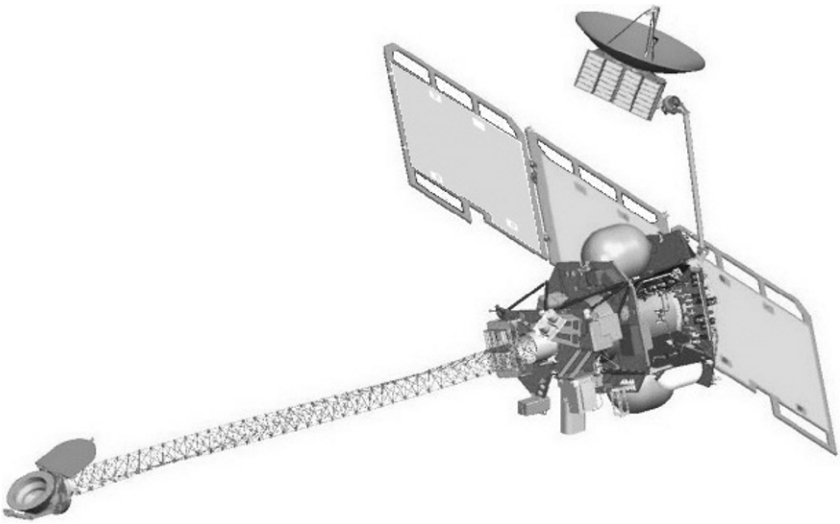

Fig. 3.15. *2001 Mars Odyssey* spacecraft has an articulated high-gain antenna, HGA (upper right) and articulated solar arrays. Image courtesy NASA/JPL-Caltech.

3.6 Scientific Experiments with AACS

Many of the engineering systems or subsystems on a spacecraft can also participate in experiments that directly provide valuable scientific data. Telecommunications radio can actively probe an atmosphere; Doppler shift, usually a tool for navigation, can be used to measure a natural object's mass. Attitude control can participate, too. *Galileo*'s AACS serendipitously made a scientific discovery unrelated to the normal business of tracking stars for attitude estimation. While *Galileo* was orbiting Jupiter, it was realized that because high-energy particles leave a trace in the scanner's data, the star scanner could actually be used as an instrument to measure the flux and energy of those particles, by calibrating and analyzing its noise signal. The data showed that the particles trapped in Jupiter's magnetic belts were predominantly less than 2 MeV electrons. Another discovery came in 2000 when the second magnitude star Delta Velorum-A was in the *Galileo* star scanner's field of view. The star drew notice by dimming below the star scanner's detection threshold for about eight hours. Subsequent analysis of the star scanner data, plus the work of amateur and professional astronomers, revealed that the star is an eclipsing binary, and the brightest one known [13]. The star's dim companion has an orbital period of 45.2 days. The eclipse, which lasts for eight hours once each orbit while the star's companion passes in front, causes a dimming that can even be seen with the unaided eye.

AACS can also help scientists investigate the density of an atmosphere on a planet under study. Accelerometers can be used in this application, if the spacecraft is intended to enter an atmosphere, as in the case of *Huygens* or *Mars Pathfinder*. The amount of atmospheric drag a spacecraft will experience depends on these factors:

$$F_{drag} = 1/2 \rho V^2 C_D A \tag{3.1}$$

where

F_{drag} = force in newtons,
ρ = atmospheric density in g/m^3
V = velocity in m/s
C_D = the spacecraft's coefficient of drag, and
A = the area of the spacecraft impacted in m^2

The spacecraft's drag coefficient should be known precisely from design and test for a spacecraft intended to enter an atmosphere, as is the area exposed to atmospheric friction. With velocity known from navigation data, telemetered measurements of force sensed via accelerometers on board can then permit solving for the unknown atmospheric density.

Even if a spacecraft is not designed for atmospheric entry, it can report on the amount of torque it experiences when flying close by a planet or other object that has gas associated with it. *Cassini* flew through the watery and gaseous geysers that erupt from Saturn's moon Enceladus. While its science instruments directly sampled the plume's constituents, the torque that AACS reported helped estimate the plume's density. *Cassini* routinely flies close enough to Saturn's largest moon, Titan, to sense the upper reaches of that moon's atmosphere. As the spacecraft flies past the 5,150 km-diameter proto-Earth-like object, AACS reports the torques felt on the spacecraft body, varying over time with altitude above Titan. Some of these targeted encounters come closer than 1,000 km to Titan's intriguing surface.

The quantity of torque on the spacecraft as it flies by Titan applies directly toward revealing Titan's atmospheric density. This torque can be estimated as:

$$\boldsymbol{R}(t) = \int_0^t \{\boldsymbol{T}_{ATMOS} + \boldsymbol{\epsilon}\}\delta t \tag{3.2}$$

where

$\boldsymbol{R}$ is the accumulated angular momentum vector. Its time-derivative denotes the per-axis body torque that AACS constantly estimates. *Cassini* reports its value in telemetry, which it computes after filtering to reduce the effect of noise.

$\boldsymbol{T}_{ATMOS}$ represents the torque contributed by the atmosphere,

ϵ is a vector quantity containing small torques that integrate to near zero, such as from gravity gradients and photon pressures.
t is time, which is indicative of altitude above Titan's surface as the spacecraft flies by.

Reference [11] describes this experiment, including how the torque values are reduced to provide atmospheric density information.

In 1993 after the *Magellan* project had completed all its prime scientific objectives at Venus and a number of extended-mission objectives, the spacecraft was also used to study Venus's atmospheric density as a function of altitude by measuring torques the atmosphere exerted on the spacecraft. The craft had two rectangular solar-photovoltaic panel appendages whose drive mechanisms could rotate them about one axis. AACS canted the two panels in opposite directions, making a "windmill" out of the spacecraft. Orbit trim maneuvers then lowered *Magellan*'s orbit periapsis, its closest point to the planet, until it was dipping into the high atmosphere. The craft's AACS reported on the RWA rotational speeds resulting from the torque to help characterize the free-molecular flow in the upper reaches of our sister planet's CO_2 atmosphere. The experiment is described in reference [14].

3.7 AACS Faults and Protection

We'll visit the subject of fault protection more specifically in a Chapter 5, but before leaving the subject of AACS we should characterize a few more of its responsibilities in regard to the basic need for reliability mentioned on page 91. AACS can take care of itself in the remote reaches of the solar system by recognizing "routine" problems as well as extraordinary ones. It does this by running software routines called *fault-protection monitors,* each of which is tasked to watch for a specific kind of problem. *Voyager*'s AACS has dozens of fault-protection monitors watching for limits to be violated or failures to occur. An advanced AACS such as *Cassini*'s has hundreds of fault-protection monitors. We considered the case of anomalous thrust, which is one of the extraordinary anomalies AACS fault-protection monitor routines look for. Additional monitors are triggered in cases such as when AACS cannot find or identify a needed celestial reference, or if it were commanded to point an instrument too close to the Sun, or when the reaction wheels are reaching their momentum saturation.

Normally, routine command sequences include reaction-wheel desaturation maneuvers at intervals that keep the wheel speeds well within limits. Should momentum build up unexpectedly in an RWA, or if regular commanding were to neglect RWA speeds inadvertently, AACS's fault-protection response algorithms would automatically interrupt the regular sequence of commands executing to perform an RWA momentum desaturation. On some spacecraft this automated step might take place routinely, and on other spacecraft it would constitute an extraordinary anomaly.

Additional fault-protection monitors can invoke built-in fault-protection response algorithms to take appropriate action in just about every kind of imaginable anomaly. Many can autonomously swap over from a failed part to a spare. And

AACS is ultimately called upon when other systems detect problems that require interrupting the normal sequence of operational commands. The request to AACS may be to rotate the spacecraft to an attitude known to be thermally safe, and which will permit communications with Earth for troubleshooting and repair.

Notes

[1]This is *Voyager 1*'s distance from Earth as of December 2008, at which time the spacecraft is on roughly the opposite side of the Sun from the planet, and heading away and north at 3.6 astronomical units per year on its hyperbolic solar-escape trajectory.

[2]In 1950 the English mathematician Alan Turing (1912–1954) proposed a test: A human engages in a natural-language typewritten conversation with a machine, which passes the test if a human judge cannot reliably tell it is not another human.

[3]The *Voyagers* are too far from Earth to use their low-gain antennas for communications. Their only choice is to point their high-gain antennas accurately.

[4]The use of reaction wheels is an alternate to direct thruster control, although these devices also require occasional use of thrusters to manage their own rotation rates.

[5]Patterns of the distant stars do not change appreciably despite a spacecraft's travels throughout the solar system. Their great distances prevent parallax from interfering with AACS's ability to recognize them in the same patterns familiar to us from Earth.

[6]The three-axis-stabilized *Voyagers* are routinely commanded to execute rotations about their Z-axes for the benefit of fields and particles investigations.

[7]The *Pioneer* missions were all managed by NASA's Ames Research Center.

[8]This remarkable animation by Dan Maas (1981–) of Maas Digital LLC, of the Mars Exploration Rover launch and mission, includes spin-up and yo-yo controlled de-spin following the upper stage burn: http://www.maasdigital.com/mervideo-large.html

[9]During design, the *Voyager* Sun sensors were modified, including the addition of amplifiers, to permit their use beyond Saturn [15].

[10]See the *Voyager* Project press release:
http://www.jpl.nasa.gov/news/features.cfm?feature=548

[11]The French physicist Leon Foucault (1819–1868) coined the word "gyroscope" in 1852 when he was attempting to use a gimbaled spinning-mass device to observe the Earth's rotation. The attempt failed due to friction and unwanted torque in his system, and Foucault is better known for his use of a pendulum to display our planet's daily motion. Any device that enables one to see rotation is worthy of the name gyroscope, whether or not the device itself involves a rotating mass.

[12]Note that inertial attitude references for a spacecraft represent a different discipline from that of inertial *navigation* in aviation and other Earth-based applications. Inertial navigation systems serve to model the vehicle's entire progression from one point to another by precisely measuring and tracking all its accelerations. While there may be accelerometers aboard an interplanetary spacecraft, they are used for tasks other than point-to-point navigation.

[13]Piezoelectric materials, typically crystals or ceramics, expand and contract in response to the application of an electric current. They also generate an electrical current when mechanically compressed or stretched. A crystal earphone demonstrates the former effect, and the latter effect is employed in the household push-button spark generator used to light a cooking flame.

[14]Also called micro-machines and micro systems technology.

[15]Coriolis effect, described in 1835 by French scientist Gaspard-Gustave Coriolis (1792–1843), is an apparent deflection from a straight path of a moving object, when viewed from a rotating frame of reference. Air masses moving south in Earth's northern hemisphere are deflected west as seen from the rotating surface, due to Coriolis effect.

[16]As of late 2008, *Voyager 1* and *Voyager 2*, launched in 1977, have used up little more than two thirds their 100-kilogram complement of propellant.

[17]Don't actually try this at home! The rapidly spinning massive wheel would pose a danger of personal injury.

[18]See http://www.spitzer.caltech.edu

[19]Again, don't actually try this, because the spinning concrete mass would present a danger of personal injury.

References

[1] Peter C. Hughes. *Spacecraft Attitude Dynamics*. Dover Books on Engineering. Dover Publications, 2004.

[2] Richard A. Kerr. Voyager 1 crosses a new frontier and may save itself from termination. *Science*, 308(5726):1237–1238, May 27 2005.

[3] W. McLaughlin and D Wolff. Voyager flight engineering: Preparing for Uranus. Technical Report AIAA-85-0287, AIAA, 1985.

[4] Walter Hohmann. *Die Erreichbarkeit der Himmelshkörper*. R. Oldenbourg, Munich and Berlin, 1925.

[5] The Royal Society of London. *On Governers*, volume 16, http://journals.royalsociety.org, 1867/1868.

[6] Richard C. Dorf and Robert H Bishop. *Modern Control Systems*. Prentice Hall, Inc., New Jersey, 10th edition, April 2004.

[7] Jean-Claude Samin. Mechanics of multibody systems. PDF on web: http://www.tele.ucl.ac.be/PEOPLE/PS/FSAB1202/FSAB1202-meca.pdf, January 2005.

[8] Alan B. Binder. *Lunar Prospector: Against All Odds*. Ken Press, Tucson, Arizona, February 2005.

[9] Cambridge Philosophical Society. *Proceedings of the Cambridge Philosophical Society*. books.google.com, 1892 (Digitized 2007).

[10] B. Dobrotin, E. A. Laumann, and D. Prelewicz. Mariner limit cycles and self-disturbance torques. Technical Report AIAA 69-844, American lnstitute of Aerononautics and Astronautics, 1969.

[11] A. Lee and G. Hanover. Cassini spacecraft attitude control flight system performance. Technical Report AIAA 2005-6269, AIAA, 2005.

[12] E. F. Koch. The linear boom actuator designed for the Galileo spacecraft. In *The 17th Aerospace Mech. Symp. p 81-96 (SEE N83-24881 14-39)*, pages 81–96, May 1983.

[13] Sebastian Otero, Paul Fieseler, and Christopher Lloyd. Delta Velorum is an eclipsing binary. IAU Information Bulletin on Variable Stars, December 2000.

[14] Christopher A Croom and Robert H Tolson. Venusian atmospheric and Magellan properties from attitude control data. http://ntrs.nasa.gov/, The George Washington University, Joint Institute for Advancement of Flight Sciences, Langley Research Center * Hampton, Virginia National Aeronautics, August 1994.

[15] David W. Swift. *Voyager Tales: Personal Views of the Grand Tour*. American Institute of Aeronautics and Asttronautics, 1997.