18 February 1983, Volume 219, Number 4586 SCIENCE

The Autonomous Viking

Edward Hutchings, Jr.

The Viking mission to Mars in 1975 is one of the brightest chapters in the history of space exploration. Despite the worldwide press and television coverage of the spectacular mission, however, there are still a few missing pages in the account. One of them is the story of how a handful of people at the Jet Propulsion Laboratory (JPL) of the California Institute of Technology kept the Viking spacecraft functioning in the last years of the 5-year-long mission.

The mission had two major goals: to study the atmosphere, geology, and soil of Mars, and to search for life in two specific locations. The overall project was managed by NASA's Langley Research Center, while JPL was responsible for developing the orbiters, tracking and data acquisition, and the mission control and computing centers.

Summary. The two Viking spacecraft launched to Mars in 1975 were designed for 90 days of intense observations followed by an extended mission phase to end in 1978. Because the spacecraft were still operating so well in 1978, three more mission phases were added and the project was not officially terminated until 1980. During these last three mission phases delays in controlling the orbiters from the earth increased. The spacecraft were kept functional and the length of the Viking mission was extended because the ground crew, over a period of 2 years, gradually made the orbiters autonomous.

Each of the two Vikings launched toward Mars was a double spaceship. One part, the orbiter, was to circle the planet continuously, photographing the surface and analyzing the atmosphere. The other part, the lander, was to settle on the surface, probe the soil, and radio back its discoveries.

Viking 1 was launched on 20 August 1975. The lander separated from the orbiter and settled down on the surface of Mars on 20 July 1976. Viking 2 was launched on 9 September 1975, and its lander touched down on Mars on 3 September 1976. Designed for 90 days of intense observations, the two spacecraft kept functioning for several years; thousands of new pictures and invaluable scientific information were obtained, providing additional insight into the Martian surface and atmospheric features.

When the primary Viking mission was concluded on 15 November 1976, an extended mission began. It was planned to end on 1 April 1978, along with Langley Research Center's management of the project. However, the Viking spacecraft were still operating so well at that time that a third mission phase, called the Viking continuation mission (VCM), was initiated. Jet Propulsion Laboratory was assigned full responsibility for the project, and the VCM was set to run until 1 February 1979.

As it turned out, two more mission phases followed the VCM, and the project was not officially terminated until 30 September 1980—successful beyond all expectations. Orbiter 2 operated until July 1978, Lander 2 until April 1980, and Orbiter 1 until August 1980. Data were obtained through more than two Martian years, and Lander 1 is still acquiring data in an automatic mode programmed to continue until 1994, barring component failure. With 52,000 photographs from the orbiters, 4500 from the landers, and vast quantities of other data, the Viking mission has provided a mine of information that scientists will be working on for years to come.

The author is a lecturer in journalism in the Division of Humanities and Social Sciences at the California Institute of Technology, Pasadena, 91125.

Two Additional Mission Phases

As originally conceived, the mission phases following the VCM were to consist of extremely low-level operations, to take limited and specific types of data, including visual imaging to supplement the information obtained in the earlier mission phases. But each of these subsequent phases produced far more information than had ever been planned, largely because of the close cooperation between the science team and the flight team. The Viking scientists, with an already superb overview of the planet Mars, kept opening up more and more interesting areas for investigation so that the project was constantly being pressed to acquire as much data as possible.

In this environment, it was the responsibility of the project manager and the project scientist to make sure that the primary objectives of the mission were not being compromised by diversionary or ad hoc requests for scientific or engineering changes that might result in excess use of attitude gas, excessive command loads, or excessive playback loads. Still, it was a common goal of the science team and the flight team to make Viking as productive as possible. Even though there has always been some division of interests between the engineers and the scientists working on spacecraft missions, the engineers on the extended Viking mission had been flying spacecraft long enough to understand their common purpose—to extend human knowledge. So they began to search through drawings, documents, and mental catalogs for the special capabilities of the spacecraft, looking for ways to help offset the austerity of the operation and improve productivity and safety.

This was not an easy task. The manpower to support their effort diminished day by day. Whereas during the primary mission there had been about 1000 people supporting Viking operations, the flight team dropped to approximately 450 for the extended mission, 170 for the VCM, then to 100, and ultimately to fewer than 25 people for the combined project and flight team in the later mission phases. These drops were occasioned not only by the reduced budgets for Viking operations, but also by the growing need at JPL for experienced personnel to work on the Voyager mission (to Jupiter and Saturn) and Galileo (to Mars and Jupiter).

The diminished Viking team included part-time workers, some newly assigned to their jobs, and some half-dozen who had developed a high level of expertise from their long involvement with Viking. These were the mainstays of an inventive, dynamic team whose collective charge was to monitor the health of the spacecraft, make sure the scientific goals for gathering data were achieved, and develop strategies to keep the spacecraft functioning. Because the team was so small, it was close-knit and could function fairly informally. There was no longer a full family of checks and balances. The operation changed from a complex one, where everything had been subjected to triple scrutiny before approval, to one where there was no one looking over anyone else's shoulder—which also meant that a slight miscalculation could mean the end of the mission.

Spacecraft's Potential for Autonomy

One of the things the Viking team recognized at once was that the spacecraft had a potential capability that was not being exploited, a fact that had not been so obvious when there were 1000 people in the operation. Not until well into the VCM phase, however, was there time to begin exploiting this capability.

The Viking orbiters each contained two programmable subsystems: a Computer Command Subsystem (CCS) and a Flight Data Subsystem (FDS). The CCS served as the central controller of the spacecraft, executing ground commands and storing sequence commands and commands related to safety. The FDS gathered low-rate engineering telemetry data and high-rate science data from the various spacecraft subsystems and sent them to the ground. During the primary mission, the orbiters did virtually nothing on their own initiative. Commands for all engineering and science operations were sent to the CCS from the earth, and the CCS saw that they were carried out. Telemetry from the FDS reported the state of the spacecraft batteries and any operational anomalies. Commands for all housekeeping operations such as charging batteries and remedying defects were sent from the earth, which involved an inherent delay in command generation and transmission time.

As these delays increased during later phases of the mission, the team began to look for ways for the spacecraft to detect and respond automatically to certain needs without ground intervention. A study of wiring diagrams of the CCS and the FDS showed that there were connections between them. Although these connections had never been used, they should allow access by the CCS to the telemetry that was being sent to the earth. This could make it possible for the CCS to command certain chores that had previously been commanded only from the earth.

There were people who doubted that this capability really existed, and others who claimed that it was never meant to be used or even that the interconnections had been cut before the spacecraft was launched just to preclude such tinkering, which might endanger the spacecraft. There was only one way to find out. The CCS engineer loaded a simple test routine into each of the orbiters, in case it might work in one but not the other. The routine was run to see whether any engineering data were coming in on the line from the FDS to the CCS.

Although the test routine was simple, the development and loading procedures were not. Before the routine could be put into the spacecraft, it had to be modeled and tested on the ground. The simulation worked, and the routine was entered into the spacecraft. Blocks of data were taken from the interface between the FDS and the CCS, put into the CCS memory, and read out from ground telemetry. The blocks of data were time-tagged so that a correlation could be made with what the spacecraft was doing at the time. In this way the data transfer capability was demonstrated.

As soon as this was done, the CCS engineer went to work on a routine that

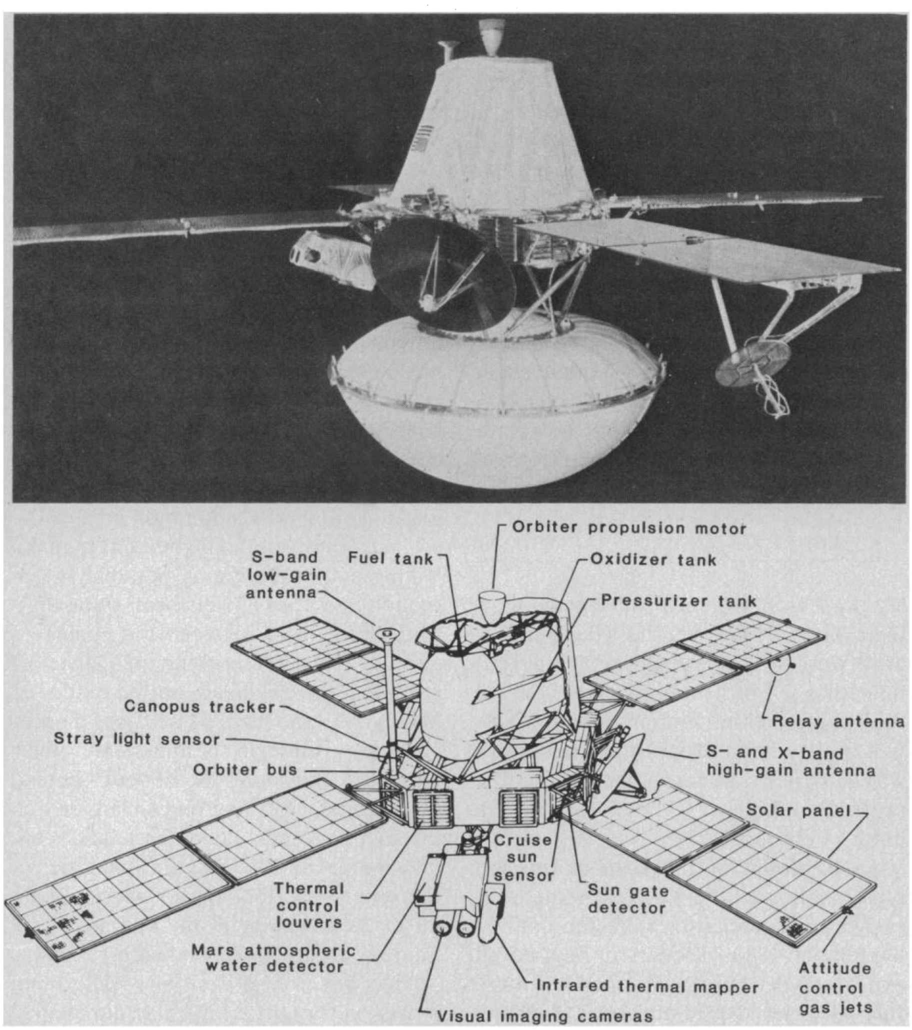

The Viking spacecraft (top) consisted of an orbiter and lander (the latter shown encapsulated in a bioshield). Details of the orbiter's equipment are shown in the diagram at the bottom. Subsystem components including the Computer Command Subsystem and the Flight Data Subsystem were mounted in the orbiter bus.

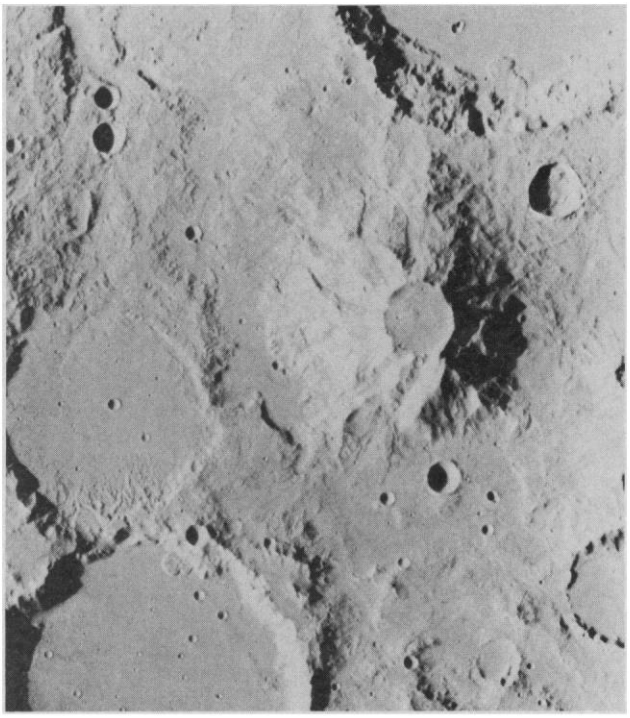

(Left) Orbiter 1 image of Mars taken in May 1980, showing a 2-kilometer-high volcano surrounded by heavily cratered terrain. The volcano is about 30 kilometers in diameter and has a summit crater 8 km in diameter. Radial channels on the flanks of the volcano mark the path of lava flow. (Right) Two-image mosaic of photos taken by Orbiter 1 in June 1980, showing ridges and grooves etched into the surface by the wind.

came to be known as DECOM (for Low Rate Engineering Decommutator Executive), which involved breaking the spacecraft data down into its component parts to isolate the data the team wanted to work with. This was a giant step. On the ground there were large sophisticated computers to do this kind of job. But the CCS was a small computer with only 4096 memory words and with only 1536 of these available for use, the rest being devoted to science and essential housekeeping routines.

The DECOM program went through a developmental stage that took several months of trial and error. This started with a program that would look at just two words, and it was first necessary to demonstrate that the computer could find these words. Next, a series of routines proved that, if the spacecraft was commanded to do a specific thing at a specific time, it could perform that function and then verify that it had done so. All this had to be firmly established before the Viking team was ready to entrust the spacecraft with any real onboard responsibility. Looking at its own data was one thing, but now it was going to be configured to make decisions predicated on what it understood it was doing.

Computers have always made decisions, but traditionally, for reliability, they have used voting techniques. Two or three computers on the line looked at the same data and some form of data coincidence provided the basis for a decision. The Viking team did not have this voting luxury, so they had to devise other appropriate checks and balances in the system. One technique was to examine the data repeatedly to make sure they fell within a particular range before they were considered valid. In other words, if a number was supposed to read in digital numbers between 1 and 127, and it read out 1 instead of the 65 the crew was expecting, the data would not be considered valid. But if the number was supposed to be 12, then anything in the range from 10 to 15 was acceptable and the data could be used.

In its final form, DECOM was able to gather requested engineering data from the FDS, send the data to the "user," and then, after confirming that the data were valid, give the user control to process them. This onboard processing of telemetry data was a capability that was new to JPL interplanetary missions. It made it possible for the CCS to automatically handle functions that were formerly commanded from the ground. And it made it possible to perform functions that were not practicable otherwise.

Onboard Detection and Correction of Gas Leaks

DECOM was not only the first necessary step in building autonomy into the Viking orbiters, it was the mechanism that made all further steps possible. It is ironic, then, that it was only by chance that DECOM got into the orbiters at all. Although it was certainly imaginative and was universally supported by the Viking team, it might never have been put to use if Orbiter 2 had not developed a serious engineering problem—a gas leak in one of the jet valves of the attitude control system that was becoming more and more difficult to deal with effectively from the ground because of its increasing frequency and because of longer delays in communication between the spacecraft and the ground crew.

The Viking's attitude control system provided three-axis (pitch, yaw, and roll) stabilization and orientation of the spacecraft. The orbiters used two celestial references for orientation: the sun and a selected star. Photoelectric sun sensors, mounted on the solar panels, locked onto the sun. With the star sensor on the star, the orbiters were kept in relatively fixed and stable attitudes. When either the sun or the star could not be seen because of occultation by Mars, inertial reference units (gyros) provided attitude references.

Attitude control jets (12 in all) were located at the tips of the four solar panels, and the attitude control system automatically commanded the opening and closing of the jet valves to release bursts of cold nitrogen gas to nudge the spacecraft into the desired orientation whenever a given amount of drift had taken place. The jets were fired several times a day.

Gas was one of the few expendable items carried on the orbiters. Each spacecraft started with about 31 pounds of attitude control gas. Inevitably, when the spacecraft used up the gas its attitude

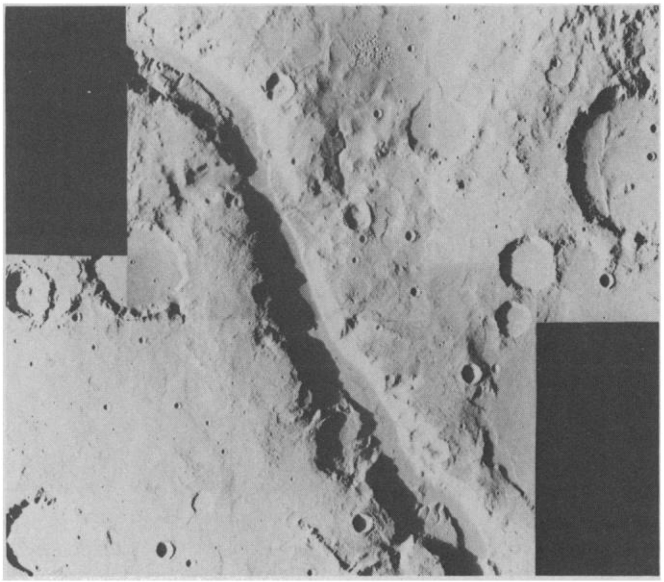

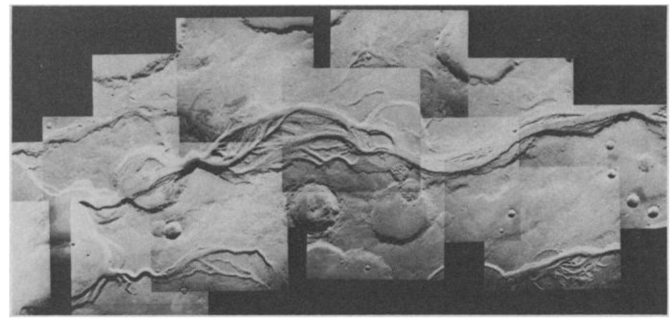

(Left) Mosaic of Orbiter 1 images taken in May 1980, showing a Martian canyon at least 800 kilometers long, about 16 kilometers wide from rim to rim, and about 2 kilometers deep. The sinuous course suggests that the canyon was produced by water erosion. (Right) Channel on Mars, probably formed by liquid water, photographed by Orbiter 1 in 1980.

could no longer be controlled and it would tumble. In the early months of orbiting Mars, the two spacecraft used gas prodigally to align themselves for science observations and to receive and relay the flood of data coming from the landers. More often than not, each orbiter was expending more than 100 millipounds of gas a day. This rate of use came to an abrupt end in the extended mission, when it was decided to go on to a continuation mission.

Historically, it was assumed that leaks were caused by tiny particles caught in the regulator valves. A deposit of particles built up in the valves after thousands of uses and kept them from being completely shut off, and gas escaped from the partly closed valves. At first, when the Orbiter 2 leak occurred, it was cleared up by opening and closing the leaking valve several times, allowing gas to flush out the contamination.

All this action was taken from the ground. When the leak was observed through the telemetry data, a command was transmitted to the spacecraft to clear the leak. A round-trip light time later (up to 42 minutes), the telemetry data would be observed again to see if any further action was needed. Even under the best of conditions, leaks could last for an hour or more before corrective action was taken at the spacecraft, and the procedure required continuous station tracking and around-the-clock staffing by ground personnel. In the VCM, neither the tracking nor the staffing was available.

Two-way communication between the Viking orbiters and the ground crew was made possible by the Deep Space Network, a worldwide tracking and data acquisition facility with antennas at Goldstone, California; Madrid, Spain; and Canberra, Australia. These stations are spaced on the earth so that one always had Mars in view, 20 million miles away, during the prime mission. By the time of the VCM, the network was sharing its time not only with Viking and Voyager, but with the Pioneer Venus orbiter and with Helios, the German flyby of the sun. Sometimes all four of these spacecraft were in the same segment of the sky, which meant that they were all in competition for visibility since the antennas in Australia or those in Spain could not be looking at all four at the same time. Therefore Viking had to compete for visibility to play back the telemetry data it had acquired, to track the spacecraft position, and to check on the health of the spacecraft.

The condition of the spacecraft was vital. As the leak on Orbiter 2 became more severe it became harder to control through ground action. Limited station coverage compounded the problem. Finally, it became clear that the problem would have to be handled by some type of onboard detection and correction. So it was out of necessity that the DECOM routine was incorporated. The routine consisted of monitoring attitude control behavior. When attitude excursions exceeded a certain limit, the jet was known to be leaking and the CCS would issue the command to turn on the leaking jet as a clearing action. The observations and clearing action were repeated until the leaking stopped.

The Orbiter 2 automatic routine reduced the gas loss from an average of 135 to less than 50 millipounds a day, even though the leaks were becoming increasingly severe. Finally the leaks could not be stopped by any routine and Orbiter 2 ran out of gas on 25 July 1978, almost 3 years after launch—its life lengthened and mission value extended to a considerable degree by the leak-clearing routine. Orbiter 1, which kept working for two more years, never developed the leaking problems of Orbiter 2. Although a six-axis leak-clearing routine was incorporated, there was never any indication that it was activated.

Monitoring the Spacecraft's Batteries

Another of the early steps in the automation of the Viking orbiters was to monitor the spacecraft's batteries. Primary power on the orbiters was supplied by four photovoltaic solar panels, which converted solar energy to electrical energy when their sensitive surfaces were facing the sun. Secondary power was provided by two rechargeable nickel-cadmium batteries, used during periods when power demands exceeded the available solar array power—particularly during occultations of the sun, which occurred daily.

By the time of the extended mission the Viking spacecraft had been exercised far beyond its original design expectations. Mechanical elements—motors, tape recorders, switches, and shutters—have finite lives. In the extended mission many of these elements, although still performing satisfactorily, had run three or four times beyond their design life. The issue became one of keeping the spacecraft going.

The orbiters were now also operating outside their design goals from an environmental point of view. The spacecraft were designed to be illuminated by the sun practically all the time, except for brief periods (about 2½ hours) of sun occultation each day. The spacecraft were actually designed with some margin to allow for 3.6 hours of occultation. Two batteries had to supply power for that period of time. But as the mission ran on beyond all design expectations, the spacecraft began to experience occultations as long as 4½ hours. The aging batteries became a prime concern. They were no longer at their full capacity, and it was necessary to preserve enough energy in them so they could survive these extended occultations. It became routine then to start turning off instruments so that the spacecraft would have enough energy to reorient after an occultation.

At first this was a ground action. But again the lack of manpower and the long delays in communications between the spacecraft and the ground crew, because of limited station coverage, made it imperative that there be some better way to keep the batteries healthy. At first the Viking team incorporated a routine to monitor the current usage of both batteries. If the telemetry showed that the batteries were delivering current unequally, the battery monitor would automatically start shutting off the instruments that were using power, whether scientific data were being collected or not, according to a built-in table of priorities. But in a long occultation there were prospects that loads other than instruments—even heaters—might have to be shutdown, which meant that there was a potential for permanent damage. To keep the orbiter functioning, the flight team had to take that chance.

Now that they knew they could monitor the battery data, the crew built a routine that allowed for optimum charging of the batteries after an occultation. This routine looked at the temperature of the battery, as opposed to the current discharge, since it was recognized that during its charge cycle the battery was endothermic, while when it was fully charged, or close to it, it became exothermic. Following an occultation, about 1 hour after the spacecraft reacquired the sun, the battery charger routine started the batteries on high-rate charge. It would monitor both the battery current and the temperature, independently, until the point of turn-off was reached. This turn-off, ensuring that the batteries would be shut off automatically at the optimal charge, was the first routine that completely took over all the work connected with battery and power that had been done on the ground by a large group of people, who had to decide on rates of charge and duration of charge through simulations by ground computer, depending on the length of each occultation.

Stray Light Detection and Protection

Another function the spacecraft took over from the ground crew was that of automatic stray light detection and protection. Stray light was the name given to random light sources that moved into the field of view of the star tracker and distracted it from its steady reference star fix. This light might come from other stars, from Mars itself or one of its moons, or even from stray particles or chips of paint flying off the orbiter.

Although these particles were extremely small, if they came close to the tracker they would often appear to have the intensity of a star. (Ground tests showed that a particle with a diameter of 0.005 inch and a diffuse reflectance of 0.4, 4 feet in front of the tracker, would be eight times brighter than the reference star Canopus. A flat, specularly reflecting particle of the same size would have the same apparent brightness 2500 feet away from the tracker.)

The routine developed to deal with stray light monitored the intensity of the reference star tracker. Whenever this intensity began increasing too rapidly—a signal that there was some kind of interference—the spacecraft would automatically switch to a roll inertial mode from its celestial mode, which meant that the star tracker was not being used, and the gyros were turned on so that the spacecraft would not drift off the reference position. The intensity of the light would continue to be monitored until it was back in the right range, at which time the orbiter would automatically go back into the celestial mode again, fixed on the star.

The routine had other features. For example, excessive stray light could damage the tracker, so at a given threshold the tracker would be shut off completely (not just the input of the tracker to the control system, as was usually done) in favor of gyro control. This meant that the light source would not be monitored. Half an hour after such a turn-off the tracker would be turned on again to check the field. If the light was still too bright, the tracker was shut off again for protection. If not, it would return to its normal tracking mode.

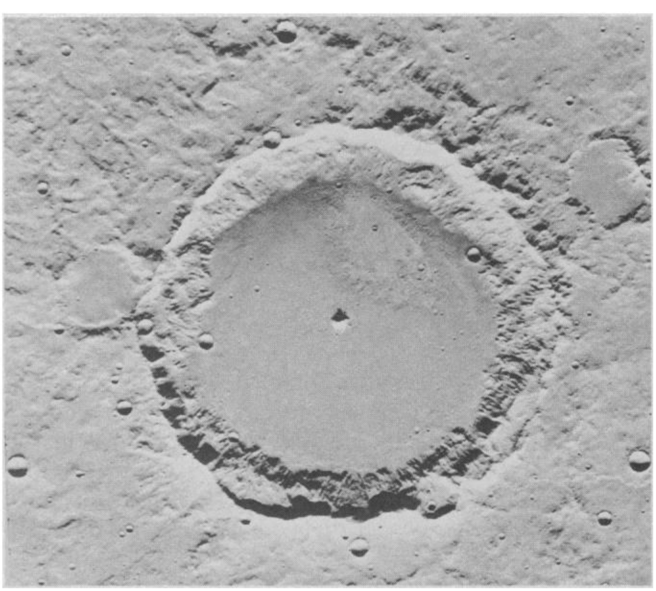

(Left) Next-to-last image taken by Orbiter 1; this is the last existing photograph of Mars from Orbiter 1 since the last image taken did not come out. (Right) Recent image from Viking Lander 1, received and processed in 1982.

Out of Gas

After orbiting Mars for 4 years, 1 month, and 3 weeks—on a mission that NASA initially hoped would last 90 days—Orbiter 1 ran out of gas on 6 August 1980. Between them, the two Viking spacecraft had provided more than 50,000 photographs, giving cartological data on 97 percent of the surface of Mars, and about 3 million weather reports. Even now, Lander 1 continues to send back daily meteorological reports and periodic photographs, and may well continue to do so until 1994.

One NASA requirement that had to be met as an element of bringing the Mars mission to an end was verification that the spacecraft transmitters were shut off, so that the Viking radio frequencies could be used again in the future without there being any chance of interference from the Viking transmitters.

By the summer of 1980, the gas supply on Orbiter 1 was running so low that it was clear the spacecraft could not operate much longer. However, by this time, in compliance with the NASA requirement, one final automatic routine had been incorporated to ensure that the transmitters would be turned off before the spacecraft went out of control. This routine made it possible for operations to continue until the gas was completely gone. It monitored the gas supply and the spacecraft attitude with respect to the sun reference. Indication of gas depletion (a gross pressure measurement) as well as loss of sun orientation (a departure of 5° from the sun line) would signal the end of the supply and the end of the mission. The routine, which also ensured that there was enough battery power to complete the operation, would start a table to shut down the transmitters. Orbiter 1 was turned off on 6 August 1980, its gas gone and its transmitters silent.

Thus the Viking orbiters, over a period of about 2 years, were made autonomous step-by-step. The results were spectacular. The spacecraft performed better, manpower and money were saved, the length of mission operations was extended for both orbiters, and more data were gathered with less effort.

The Viking orbiters were probably the first spacecraft to incorporate and depend upon this degree of autonomy, but there will be others. Autonomy in spacecraft systems will be a design goal of the future. But it is not likely that there will be another autonomous spacecraft produced out of such need, by such a remarkable crew and such a small one. And surely there will not be another in which the autonomy is incorporated while the spacecraft is circling a distant planet.

Actions of Estrogens and Progestins on Nerve Cells

Donald W. Pfaff and Bruce S. McEwen

Summary. Estrogens and progestins alter electrical and chemical features of nerve cells, particularly in hypothalamus. Temporally, these events follow nuclear receptor occupation by these steroids, although not all effects have been proved to depend on translocation of receptors to the nucleus. Narrowing studies to focus on particular medial hypothalamic cells has been useful for understanding some of the actions of these steroids in brain. The variety of morphological, chemical, and electrical effects allow for a multiplicity in the cellular functions controlled by these hormones.

In addition to responding rapidly to electrical signals from other neurons, nerve cells react more slowly and over longer periods of time to various chemical messengers arising from within the brain and from without. Among the most potent and specific of chemical signals from outside the brain are the steroid hormones, long recognized as regulators of patterns of behavior related to reproduction and defense of territory (*1*) and more recently recognized as having influences on mood and affective state (*2*).

The cellular mechanisms of steroid hormone action have gained attention in the past decade through the discovery of specific receptor-like macromolecules in the soluble fraction of target tissues which bind the hormone and carry it to the cell nuclear compartment (*3*). This has led, in turn, to incisive studies of the consequences of the actions of hormones like estradiol with respect to how new protein synthesis is induced in the uterus and its relation to cell growth, cell division, and other events (*4*). Similarly, from this model, the manner in which estradiol and progesterone interact with receptor macromolecules in the chick oviduct, trigger messenger RNA (mRNA) formation and, in turn, synthesis and secretion of specific proteins, could be studied in such detail as to approach a molecular biology of hormone action (*5*). A frequent theme in such studies is that hormone occupation of receptor sites in the cell nucleus, for more than a few minutes and often for several hours, is required for the full range of hormone effects (*6*).

The brain is no exception as far as steroids are concerned and has been found to contain receptors for all five classes of steroid hormones: estrogens, progestins, androgens, glucocorticoids, and mineralocorticoids (*7*). These receptor sites are not uniformly distributed but rather occupy specific loci within the brain. As a consequence, studies of the localization of specific steroid hormone–concentrating nerve cells are especially important in brain, where small regions are involved in various complex functions. Specialized steroid hormone autoradiographic techniques (*8*) have the advantage of detecting hormone-accumulating neurons amidst other types of cells. This technique has been used for studying estrogen and androgen receptors among vertebrate classes of fish, amphibia, reptiles, birds, and mammals (including rodents, carnivores, and a primate—the rhesus monkey) (*9*). The neuroanatomical lawfulness of sex hormone binding in the brain is considerable. All species studied have estrogen- or andro-

Dr. Pfaff is professor of neurobiology and behavior and Dr. McEwen is professor of neuroendocrinology at the Rockefeller University, 66th Street and York Avenue, New York 10021.