THE IMPACT OF ROBOTS ON PLANETARY MISSION OPERATIONS[1]

A. A. HOOKE, B. T. LARMAN, and W. M. WHITNEY
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Summary. For reasons of efficiency and safety, unmanned roving vehicles sent to explore remote planetary surfaces must carry out some of their tasks without step-by-step human control. To realize the benefits that such semiautonomous machines can provide will require some changes in how planetary missions axe presently planned and conducted. Specifically, mission profiles will have to be based on tasks or functions rather than sequences of timed events, scientists will have to be more directly involved in the control of their instruments, and present ideas concerning spacecraft safety, testing and simulation of vehicle performance, telemetry design, and ground-system implementation must be reexamined.

Introduction. The families of spacecraft presently used in planetary investigations have limited ability to cope with unforeseen problems. For the orbital and flyby missions conducted to date, the return of scientific data has been bountiful and the spacecraft have been highly successful despite inflight anomalies, because most of the external and internal conditions with which they have had to contend have been anticipated and taken into account in design. Other kinds of missions, however, especially those involving exploration of the surfaces of planets or their satellites by roving vehicles, will require that the remote machines operate in environments whose detailed features cannot be known in advance. Survival, and the efficient execution of tasks, will demand that the rovers be able to do for themselves some things that would otherwise require protracted step-by-step control by people on Earth.

The purpose of this paper is to examine how the transfer of more autonomy to spacecraft will affect the manner in which planetary missions are planned and conducted. The concern is not with technical problems in artificial intelligence and robotics, but rather with some of their consequences when solutions axe found. The spacecraft is only one element in a complex chain of systems needed to conduct a mission. The allocation of functions and responsibilities between the spacecraft and the ground reflects such technical

---

[1] This work was supported by the National Aeronautics and Space Administration under Contract No. NAS 7-100.

considerations as the mission objectives, available technology, and the limitations of the communication link. There are also important non-technical considerations: the investment in money, time, and experience in capital equipment and systems that already exist; the delay and expense in making changes; an(! the dynamics of the interaction among the people who design spacecraft, those who use them as experimental tools for scientific research, and those responsible for prudent inflight operations. To ignore these factors in fashioning solutions to the technical problems is to guarantee that the solutions will not be used.

In this paper we describe briefly how present missions are planned and how they are carried out. The need for more spacecraft autonomy is reviewed for a Mars roving vehicle of the kind that might be flown in the period after 1985. Possible scientific objectives, functional requirements, and characteristics of such a rover are given and a scenario for its use is developed. We call attention to aspects of our present mission-control philosophy and approach that will not be adequate for more autonomous spacecraft, and suggest some changes that will have to be made.

Mission Design. Although planetary spacecraft have evolved into rather complex systems, the manner in which they are presently flown remains, at least in principle, relatively simple. Mission profiles are based on pre-defined blocks of sequence events. Each block specifies related actions or changes in state of one or more of the spacecraft subsystems, and contains command components which follow a rigid timeline. Inflight sequences are built by linking the blocks and inserting or modifying commands as required. Such sequences can be executed in response to commands directly sent from Earth, or can be stored in the memory of the spacecraft computer. The computer keeps track of mission-elapsed time by counting an on-board clock signal. Critical sequences may be initiated automatically when the predetermined moment arrives on the spacecraft clock, unless a prior command from Earth to inhibit the sequence has been received. By putting blocks of sequences together in the right order, and scheduling their execution at the right times, the entire mission is constructed. Once initiated, the profile minutely follows the preplanned timeline, and will deviate from it only in the event of on-board failure or upon intervention by ground control.

The success of the above approach reflects three important characteristics of present missions. First, the times when the spacecraft will arrive at given points of its trajectory can be very accurately predicted. The spacecraft can be placed on a predetermined flight path by its launch vehicle and kept on course by corrective firings of its own propulsive system. Second, in the missions conducted to date, planetary spacecraft have not disturbed or interacted with their environment in such a way that their own subsequent behavior was affected unpredictably. Finally, despite the fact that unplanned as well as planned events occur, and that detailed scientific strategies are sometimes changed after launch, there has usually been enough flexibility and functional redundancy built into the spacecraft and enough inactive cruising time built into the flight plan to enable new sequences to be designed that work around an anomaly and achieve some or all of the mission objectives. Thus, although the sequence-block approach to mission design appears to be rather inflexible, it has in fact proved quite adaptable to increasingly complex mission profiles, especially where human ingenuity has modified the blocks in response to anomalies or changing requirements.

One can, however, envision classes of missions in which the design philosophy and control strategies that have worked so well will no longer be adequate. In the remote exploration of planetary or satellite surfaces, which is the concern in this paper, the times and the ordering of significant mission events and the interaction of the spacecraft with its environment cannot be predicted. Highly detailed preplanning of mission sequences will therefore be inappropriate, although, as will be seen, certain features of block organization can be retained.

Present Space Vehicles and Their Control. We shall describe here the procedures that are used to control our present Mariner series of deep-space probes and some of the problems faced in implementing them. Similar techniques will be employed on other interplanetary missions within the next decade, including two stationary scientific stations that are scheduled to be landed on the surface of Mars in 1976 as part of the Viking mission.

The objectives of flyby, orbital and landed missions embrace investigation of particular planetary environments and exploration of surface topography, morphology and physical properties. The probes are characterized by a large set of complementary payload instruments, supported hospitably by “engineering” subsystems of the spacecraft. Certain of the non-imaging instruments are rigidly fixed to the body of the probe, but targetable experiments such as TV cameras and spectrometers are usually co-located on a scanning platform. The probe is stabilized with respect to some celestial or surface reference, and the platform is pointed at objects of interest via two-degree-of-freedom actuators.

The spacecraft communicates with the ground system by formatting scientific and engineering data into digital telemetry streams, and transmitting these data to a receiving station on Earth. on the ground, the incoming data streams are relayed to a central processing site for decommutation, display and interpretation. On the basis of the analysis of these incoming data, commands may be generated and transmitted which initiate, modify, refine or otherwise perturb upcoming preplanned sequences. A simplified block diagram of this human-interactive feedback system is presented in Fig. 1. Specific elements of this control process will now be discussed in more detail.

Phenomena of the planetary target stimulate the science instruments, which produce electrical analogs of the disturbance. These signals are sampled, digitized into binary measurement words, and fed to a telemetry formatter. Under control of hardware or (increasingly) software, the measurements are multiplexed into predetermined formats to construct a ccimnutated data stream. on-board measurements from the spacecraft engineering subsystems are multiplexed with the science data, and synchronization identifiers are added to form composite pulse-code-modulated telemetry frames. These bit streams are then phase-modulated onto subcarriers, which in turn modulate an S- or X-band carrier in the radio transmitter. The signal is radiated from the spacecraft antenna and passes through space to a receiving antenna of the ground-station network. Since the space link may introduce considerable noise into spacecraft data, the stream is often coded prior to transmission to lessen the effect of such errors. At the ground station the signal is received, detected from the carrier and subcarriers, and decoded if necessary. The “raw” spacecraft data are inserted into transmission blocks by the on-site telemetry processor, and are then relayed in this form to the mission-control center via land and satellite links through a global data-transmission network. (It should be noted that the capacity of this network is often much less than the rate of incoming data; thus on-site editing is required to maintain the flow of information for analysis.) At the central processing facility the incoming bit streams are de-blocked and sent to a computer for frame synchronization, decommutation and display. Data are here converted back into analogs and presented to the science experimenter or the spacecraft analyst in real or non-real time for his examination. on the basis of an interpretation of these incoming data, which describe the phenomena being observed by the probe, sequence modifications may be recommended which improve the acquisition of future data or which lessen an environmental hazard. Such modifications are discussed and approved, mission sequence changes are generated by ground controllers and fed to large-scale computer programs which assemble the required commands into a file suitable for transmission back to the remote ground station. Before such transmission is approved by the mission director, the receipt of the command file is simulated by complicated software packages which model the spacecraft response. This simulation is validated by operational personnel to verify that mission constraints are not violated by the revised sequence.

The sequence generation and command validation process is extremely laborious and requires much manpower and time, since it is at this stage that all of the spacecraft and ground idiosyncracies coagulate and influence the production of a working mission profile. The task is minutely detailed, requires extreme precision, and involves many iterations to complete a flightworthy profile that reflects competing mission requirements. The process is in many ways similar to the painting of a picture, where an abstract concept is translated into colored images on canvas. Not only must the artist crystallize a large number of ideas into a composed scene, but he must also prepare the canvas, mix the paint, build the scene with a variety of different brushes, over-paint areas which are not in perspective, critically inspect the finished result, and finally sign, frame and take it to the gallery for sale. The flight-sequence-generation process has historically been a major bottleneck in the progress of a mission, and will remain so until spacecraft are able to look after themselves without large amounts of care and feeding from the ground.

Once generated and validated, the command files are reviewed and approved by the mission director, and are then sent via the global. communications network to a remote ground station, where they modulate a subcarrier applied to an S-band carrier and are radiated to the spacecraft. The signal is received at the spacecraft and the command messages are detected, decoded and fed to an on-board computer. The processed commands may modify a previously stored sequence program within the computer, or occasionally may be routed directly to a user subsystem for immediate action. The spacecraft then executes the modified sequence, thus altering data-acquisition parameters and providing new information for analysis.

Of particular concern to this discussion is the nature of the processes involved within the mission-control loop, for some of these must be synthesized by any autonomous probe. There are four major and largely separate mission-control activities. These are shown organizationally in Fig. 2 and include: (1) Monitoring of engineering parameters telemetered from the probe, and verification of correct operation. This work requires large quantities of real-time engineering data to be presented to analysts, who are responsible for maintaining continued health of the spacecraft. From their scrutiny of the performance of those subsystems which support the science instruments, the spacecraft-team engineers may recommend sequence changes which improve the quality of science return or -which are necessary to prevent immediate or future hazards which would jeopardize mission progress. Over the long term, these analysts assemble a wide data base for performance review and predictions. (2) Interpretation of science data, sometimes in real time but usually over extended intervals. on the basis of these analyses, the science team may recommend retargetting or resequencing the probe’s instruments to improve monitoring of the observed phenomena. (3) Sequence generation, based largely on a previously developed time-line, but which must accommodate the adaptive inputs from (1) and (2) above. As noted, this function entails many software operations in order to assemble, validate and integrate the new sequence. Since reprogramming of the on-board computer is usually required, a large amount of time must be expended scrutinizing a software simulation of the command series before transmission, in order not only to verify that the changes have been incorporated correctly, but to check that the modifications have not impacted or invalidated ongoing sequences. (4) Sequence execution, usually by a multi-shift team of mission controllers, led by a mission director. This group is responsible for all real-time operations of the probe and is empowered to take emergency command actions if required.

Since the components of the mission-control feedback loop are predominantly human, there are some weaknesses in the system. Firstly, and obviously, human mistakes are made: data may be interpreted incorrectly, poor conclusions may be drawn, and sequence errors may occur. Secondly, because of the limited "intelligence" of the probe subsystems, human beings must agonize over failure modes and worry about constraint violations and external disturbances to spacecraft equilibrium. Subjective judgment and emotions are involved in their decisions, which often must be made under pressure. Basic conservatism often forbids some actions which would otherwise improve science return, and adds unnecessary activity. Thirdly, people become fatigued and are effective for only limited periods of time, making it difficult to cover extended operations with key personnel. All of these factors, and many others,[1] slow down the operation of the mission, add organizational checks and balances, and may introduce frustrating delays between the scientist and control of his instrument. It is not unusual for a fairly simple sequence modification to take from twelve hours to one week to accomplish.

There are obstacles of other kinds that the mission director has to surmount in meeting his objectives. The elements of the system with which he is working may be far from optimum for their purposes.[1] The capabilities of the ground system (or lack thereof) reflect inherited decisions made for earlier missions on the basis of requirements and assumptions that may no longer be valid. Many of the spacecraft subsystems may be hand-me-downs, in design or in actual hardware, representing the philosophy of "minimum change from the previous mission." This conservatism and the money that it saves may have enabled the mission to be sold, but, as is often discovered, the money necessary to take care of an heirloom may exceed its worth.

A problem of a different kind arises from the fact that the deep-space network, the ground-control facilities, and even the mission command and control complex, axe commonly occupied in supporting not one but several dissimilar spacecraft. Consequently, in the implementation of mission-control procedures, it is not always possible to assume that these facilities will be dedicated and available when they are needed. Very often, for instance, heavy software use by one project will completely backlog another. The mission director thus finds himself constrained by outside forces beyond his control, and requiring high-level executive action to resolve priorities. Particularly severe is the problem of requiring twenty-four-hour coverage by the deep-space network to monitor spacecraft health and provide immediate command recovery should an anomaly occur. If spacecraft can be made more self-reliant, then the ground systems can be relieved of their nursery duties and can be better deployed to return scientific information.

Semiautonomous Control of Spacecraft. Remote operations similar to those that will be needed for exploration of the Martian surface have already been carried out on the Moon. Two Surveyor spacecraft had electro-mechanical scoops that were used to dig trenches, transport solid samples from one place to another, and (on one occasion), to deploy a scientific instrument and later move it to a new spot. The Russian Lunakhod vehicles were remotely driven from one site to another to gather data. These machines were adoptively commanded by operators on Earth who studied television pictures sent from the spacecraft to verify that previous commands had been correctly executed, and to plan the next moves.

At lunar distances, with a three-second round-trip communication delay and a high-rate data link for picture transmission, placing a man directly in the control loop is a slow but acceptable way of carrying out tasks with a remote machine. For a spacecraft on the surface of Mars or a more distant planet or satellite, however, this method would be suitable only for the simplest of tasks, those not requiring many individual steps or an accommodation to detailed features of the surroundings. It would also not be suitable in an emergency in which the spacecraft had to act rapidly to escape from a dangerous situation. The round-trip communication time between Earth and Mars ranges from 6 to 44 minutes. The data rate for picture transmission could vary from several thousand bits per second to one hundred times less under unfavorable conditions. Because of the rotation of Mars, the time available for surface operations would be only half of each Martian day, assuming world-wide monitoring of the spacecraft from Earth, continuous availability of the ground system, and rapid decisions by Mission Control. This time could be reduced even further if it were necessary to acquire data from the roveX through a relay satellite in orbit around Mars. In an earlier study,[2] it was estimated that, for a round-trip communication time of 27 minutes, a ground-controlled Martian rover could cover at most 300 meters during a 13-hour mission-operations period. Approximately 1 1/2 hours would be required for each 50-meter step--6 minutes for vehicle motion, the rest of the time for necessary verification and commanding operations on the ground and for transmission delay. Extensive surface investigations in this mode would be both protracted and costly.

To increase his efficiency in exploring Mars by remote control, it will be necessary for man to remove himself from some of the control loops, especially for tasks whose execution would require the transmission of large quantities of environmental data, including pictures. The machine itself will have to extract needed information from this data and use it to control its own actions.

The prime candidates for automation are the functions of locomotion and manipulation.[2,3] Ferrell and Sheridan[4] have introduced the notion of supervisory control for the remote operation of manipulators when the communication link has limited bandwidth and a time delay longer than a few seconds. Man turns over authority to the machine to complete its manipulator tasks, using data acquired from the local environment, without requiring step-by-step human control. He does, however, retain responsibility for deciding what tasks are to be done and initiating their execution. This concept is of course not limited to manipulation, and is readily extended to include locomotion as well.

At JPL, in a robotics research program sponsored by NASA's Office of Aeronautics and Space Technology, an effort is being made to demonstrate the feasibility of making a "semiautonomous robot" that will be able (on command) to move safely from one place to another, or to pick up a rock or other object, without human intervention. For these two tasks, the rover will close feedback loops that would otherwise require human interpretation of large quantities of sensory information and human decisions and commands. In other respects, the robot will be controlled from Earth, as were Surveyor and Lunakhod. Scientists and mission-control personnel on the ground will determine the scientific objectives and strategies of the mission, select the targets (including the individual rocks of interest) at each site, and plan the sequence architecture that will be required to carry out surface operations. However, those supervisory elements of the resultant command strings involving manipulation or locomotion will be decoded and executed autonomously by the machine. By combining these two capabilities for autonomous operation with other more traditional spacecraft functions controlled from Earth, we believe that it will be possible to increase the return of scientific data from a Mars surface mission by a factor of 100 over that obtained from a mission conducted entirely by direct control.

Rover Functions and Subsystems. The scientific objectives of Mars surface exploration will include investigations of atmospheric composition and meteorological conditions, surface geological features, soil and rock composition and texture, magnetic fields, local gravity, seismicity, thermal properties, and the presence of organic or living matter.[5] During the rover's design lifetime, perhaps one Martian year or two Earth years, it would be expected to gather data from many different sites, and cover a total distance of several hundred kilometers.[3,5]

To carry out the above studies, the rover would have to have a variety of scientific instruments that we need not discuss further here. In addition, it must be able to pick up and manipulate rocks and soil samples, and to deploy seismic and meteorological packages or other scientific instruments. It will be necessary for the rover to grasp and use simple tools to conduct sequences of standard laboratory operations, either automatically or in response to commands from Earth.[5] The rover will be required to move from one place to another safely and without extensive human intervention by selecting routes, guiding itself, avoiding obstacles, and determining its own position on the Martian surface accurately enough for the purposes of navigation, using stored terrain information gained from prior satellite surveys or from the rover's own studies of its surroundings.

The rover must be able to maintain or periodically to establish contact with Earth by direct radio link, through an overhead satellite relay, or via the lander from which the rover was deployed. Its data system must be able to accept and respond to commands from Earth, to manage and execute all spacecraft operational and scientific tasks, and to maintain an outgoing stream of telemetry data containing pertinent information about vehicle status and its surroundings as well as the scientific data being acquired. Execution of these functions will require provision of a large on-board base of stored data and procedural routines.

The rover will most likely be a four- or six-wheeled vehicle designed for high mobility in sandy, rocky terrain. The navigation system will, probably be built around a gyro compass, odometers, and inclinameters.[2,6] The vision system will consist of dual cameras for stereo color TV and a scanning laser range finder. Each of these instruments can give range and contour information, and can serve as backup if the other should fail. The vision system will also supply primary mission science data in the form of exploratory TV panoramas and targetted frames.

The rover will carry at least one electromechanical manipulator with several degrees of freedom and several detachable end effectors (or hands). Another manipulator--perhaps less complex--may be available to serve as a vise or holder. Each manipulator may contain tactile or proximity sensors to assist in approaching and grasping objects. Similar sensors may be placed elsewhere on the vehicle for protection from obstacles not detected by the vision system.

The vehicle will communicate by radio with its relay station or with Earth via a steerable S- or X-band antenna. This antenna and its radio subsystem may also be used for ranging and obstacle detection if the stereo TV or the laser ranging device should fail, or if there should be a dust storm that only the microwave signals can penetrate.

Scenario for Rover Operations. The following scenario represents a typical set of operational activities for a semiautonomous roving vehicle like that described in the preceding section.[2,5] After several days of terrain survey in Mars orbit, the rover will be transported to the surface and then will descend from the lander stage to the ground, a small step for a robot but a giant step for robotkind. Housekeeping time will be devoted to putting all of the vehicle subsystems into operation and checking them out, and to insuring that the communication link is well established. A panoramic view of the surroundings, including some range information for prominent features, will be acquired and transmitted to Earth for study so that the first traverse can be planned by the flight-project teams. In the meantime, scientific study of the area immediately adjacent to the landing site will begin. Stereo color pictures of surface topography will be obtained. Analyses of soil, rock, and atmospheric composition will be carried out, and meteorological and seismological instrument packages may be deployed.

When activation is complete, the rover will be instructed to go to a new site. This move and succeeding ones will be accomplished in short steps until confidence An the ability of the rover to negotiate the terrain has been established.[2] When the vehicle has reached its destination it will spend the next several days conducting a scientific study similar to but perhaps more extensive than the first. After as much information as desired has been obtained here, the rover will be directed to move on once again. occasionally, while it is in the midst of a manipulation or automation task, the rover may be brought to a halt by an inability to make a suitable plan or to execute it. It will send an alarm to Earth, and then automatically, or on command, supply pictures of the surroundings or other status information so that the problem can be diagnosed and solved by mission control.

The Impact of Semiautonomy on Mission Operations. Several characteristics of the preceding scenario have significant implications for the design of semiautonomous rovers and the development of effective mission-control procedures. We shall concentrate our attention here on three only, and then list some of their consequences.

The first characteristic has already been mentioned and is most obvious, although its importance is not diminished thereby: the exact sequence of operations to be carried out at each site, like the conditions to be encountered there, will not be predictable in detail. What is done will depend upon what is found that is judged by scientists on Earth to be of special interest, upon mission priorities, and also upon the difficulties that the machine has to overcome in carrying out its tasks. As a consequence, it will not be possible to prepare in advance an accurate timetable of rover events. Crossing a desert-like region may take one day; covering the same distance over very rocky terrain may well require a week. Exactly where and when difficulties will arise cannot be foreseen. These comments hold, incidentally, whether the mission is conducted semiautonomously or by direct control.

Second, the rover will be subjected to more hazards than orbital and flyby spacecraft or stationary landed probes. Because of its mobility, it can come to grief by bumping into rocks, tipping over on uneven terrain, or driving into a crater or off a cliff. Reckless probing with its manipulator can destroy sensitive scientific instruments or damage other parts of its “body.” Because of the communication-time delay, ground-control personnel will be powerless to stop some of these harmful movements once they are initiated. Thus, the safety of the robot will be of special concern to the mission director and to the spacecraft designers.

Third, to carry out its tasks, the rover will have to acquire a wealth of data from its surroundings, including three-dimensional topographic information. To send all of this data back to Earth for the purpose of a bit-by-bit predictive simulation of spacecraft performance and behavior will not be feasible, and would in fact negate the advantages of the semiautonomous control mode. Thus, certain information gathered by the rover and used in planning and execution of its tasks will not be available on the ground.

From the above characteristics, we are led to several conclusions about how present mission-planning philosophy and operational procedures will have to change to accommodate planetary surface exploration with a semiautonomous rover.

Mission Profile

As a consequence of the unpredictability of mission events, the planning of surface mission profiles will have to shift from a time-oriented to a task- or function-oriented basis.[6] The block method described in an earlier section will not have to be abandoned altogether. Certainly the mission can be broken down into some major segments, such as checkout, site exploration, traverse, and hibernation, and these can be further analyzed into independent functional units or blocks, as before. The blocks, organized in series to carry out a major mission segment, can be initiated and controlled by supervisory commands from Earth. The operations within each block, however, will have a dependent or conditional nature; each element of the exploratory profile will be triggered by successful completion of the preceding task, or by the satisfaction of a set of requirements or preconditions.

Scientist Involvement

The type of mission under consideration will have more the characteristics of Skylab experiments or of the Surveyor, Lunakhod, and Apollo studies of the Moon than of previous studies of the near planets from close approach or from orbit. Rover activities will be guided more by a work plan and by time “guidelines” than by a rigid timetable. As a consequence, there will be a need for much closer interactions between the scientists and their remote instruments.[6] Scientists are going to have to be more intimately involved in the supervisory control loop. It will be more important than for orbital or flyby missions for each investigator to receive data for his inspection while measurements with his instrument are in progress. When changes in the exploratory sequence seem advisable, members of the scientific and project teams will have to come to agreement rapidly about what is to be done and when. A one-week delay in instituting changes will certainly not be acceptable.

The recent trend has been away from more direct involvement of the scientist in mission operations. As mission-control procedures have grown more complex and more highly structured, the working relationships on the ground have became more formalized, involved, intricate, and sluggish. The scientist has felt increasingly divorced from control of his experiment. He will have to be put back in the driver’s seat-if not as pilot, then at least as co-pilot.

Spacecraft Safety

It seems inevitable that, just as it is necessary to transfer more autonomy to the rover for conducting its tasks, it will be necessary to give it more responsibility for its own safety and, at the same time, to relax somewhat the protective philosophy that underlies some of our present flight-control procedures. The rover must, as a first requirement, have the ability to avoid or preserve itself by reflex action from the threats presented by its surroundings or by its own “careless” actions. It must also be structured in such a way that it can protect itself from unintentionally harmful commands sent from Earth. Much of the sluggishness of mission operations results from the need to check all mission sequences for mistakes that would place the spacecraft in jeopardy. The rover will have to be designed in such a way that people on the ground can make errors in mission planning and in spacecraft operation without causing catastrophic failure. Rules for accepting or rejecting commands must be built into the machine’s operating system so that it can decide, on the basis of information about its environment and its own internal states not available on Earth, whether or not execution of the commands should be attempted. If execution is initiated but leads to trouble, the robot should be able to recognize the unwholesome concourse of events, halt its activities, make itself secure, and call for help.

Ground Simulation and System Testing

To avoid command errors that would violate on-board constraints, spacecraft performance has in the past been simulated in minute detail on the ground before a block of sequences has been transmitted for execution. For a semiautonomous rover, bit-by-bit simulation will not be possible. Functional simulation, however, will be both feasible and useful. As mentioned above, it will still be possible to plan a mission profile on the basis of functional blocks. Individual blocks can be validated and checked out on the spacecraft before launch. During the mission, functional simulation can be carried out in parellel for the purpose of predictive display and approximate task-completion forecasting, using highly compressed data relayed from the spacecraft (for example, rock sizes and positions) that the robot will use in planning its own actions. Such simulation will also grossly check for violations of major system constraints.

To insure that the rover will respond in a predictable and safe manner to the planetary environment will require a comprehensive system-testing program. How to provide an adequate environment in which these tests can be conducted, and how to structure and carry out, the testing of such a complex machine, pose major problems.

Telemetry

Up to now, the design of telemetry systems for planetary spacecraft has been able to take advantage of the predictability of mission events. It has been possible to determine the sampling requirements for data, to assign telemetry channels in advance, and thereby to hard-wire the data acquisition formatting, and commutation circuitry. Inflight changes have been restricted to commanded selection of a limited number of mode options. Only recently, with the advent of a semiprogrammable data system aboard the Mariner Venus-Mercury spacecraft, has it been possible to make a small number of modifications to the actual format of the telemetry stream after launch.

For a semiautonomous rover, however, the variable nature of its tasks, the lack of an event timetable, and the need to look after its own welfare will call for an adaptable telemetry structure.

During normal operations, it will be difficult or impossible to specify an optimum fixed-structure for the telemetry stream in advance. It may be necessary that the spacecraft should select from a variety of methods of collecting data, identifying them, and transmitting them to the ground. For example, to conserve link bandwidth, the majority of engineering measurements may only be transmitted when they violate an alarm test.

In an emergency, the spacecraft must sense internal failures or external anomalies and, in addition to coping with them autonomously, supply adequate diagnostic information to Earth. The spacecraft should be able to exercise available redundancy or corrective capability without requiring intervention from the ground, and with the minimum of disruption to on-going activities. However, enough status information must be transmitted to the controllers to enable the disturbance to be modelled and its impact on future sequences to be assessed. Much data of use in operations analysis will be present in the science data stream. To provide other environmental information, the spacecraft will have to alter telemetry formats adoptively, perhaps to increase the sampling rate of pertinent measurements, to jettison low-priority tasks, and to edit or compress other types of data that are sent to Earth. Such functions will require the incorporation of a diagnostic executive program in the rover software system. Clearly, this system must be very closely integrated with ground data-processing software; otherwise interpretive chaos will result when the robot reacts to complex stimuli. only in the event of severe hardware malfunction or unusual environmental trauma will it be necessary to override the system and resume hands-on control.

Ground-System Automation

The requirements for the ground to respond rapidly to diagnostic communications from the rover executive indicate the need for comprehensive end-to-end system design. This has not historically been a characteristic of interplanetary missions, where the spacecraft and ground-system designs have converged mainly at the interface. The autonomous rover will need very close interaction with its ground system, particularly the soft-ware, and will require a high degree of automation associated with the mission-control programs which provide sequence generation, command assembly, constraint validation, and telemetry-monitoring functions. By designing the rover and ground systems concurrently, problems and incompatibilities will be identified early and will not lurk hidden, waiting for discovery at a critical mission phase. If the design can install the scientist in a pilot-function, then mission-control response time will be reduced, and an improved dialog will occur between scientists, spacecraft engineers and the mission director. If automation makes possible a reduction in the number of mission personnel engaged in unrewarding tasks, an additional benefit may be to reduce the cost of -post-launch ground operations.

Epilogue. Loosening control over a remote machine with some autonomou.s capabilities means giving up detailed prediction of its activities. Assigning it more responsibility to carry out its tasks without human intervention lays on it the added burden of looking after its own welfare. Properly applied, these features should expedite the activities of mission control.

In all previous missions with less autonomous spacecraft, anomalies or discoveries have necessitated rapid changes in plans and procedures. By dint of enormous amounts of effort and ingenuity it has always been possible to accommodate them. While the use of ad hoc procedures is therefore not new, they have been applied to extraordinary situations only and are not features of the more conservative day-to-day operations.

For missions in which spacecraft have more autonomy, however, it will be possible to give the “fly-by-the-seat-of-the-pants” philosophy a more central position in our guidelines for spacecraft and mission design. Unless we change the manner in which spacecraft are used, we may not gain the advantages that their greater capabilities and independence promise.

References.

1. L. D. Jaffe, “Problems of Previous Projects on Scientist/machine/ Operations Interfaces,” Document 760-72, Jet Propulsion Laboratory, Pasadena, California; July 1972.*

2. J. W. Moore, G. K. Hornbrook, W. S. McDonald, J. R. Gilder, W. E. Dorroh, M. Swerdling, H. Bank, R. E. Imus, T. Gottlieb, L. Y. Lim, D. W. Kurz, and P. H. Roberts, “An Exploratory Investigation of a 1979 Mars Roving Vehicle Mission,” Document 760-58, Jet Propulsion Laboratory, Pasadena, California; December 1970.*

3. E. Heer, “Summary Report: Advanced Teleoperator/Robot Studies for Planetary Surface Roving Vehicles,”Document 760-75, jet Propulsion Laboratory, Pasadena, California; July 1972.*

4. W. R. Ferrell and T. B. Sheridan, “Supervisory Control of Remote Manipulation,” IEEE Spectrum, Vol. 4, No. 10, pp. 81-88; October 1967.

5. R. Choate and L. D. Jaffe, “Science Aspects of a Remotely Controlled Mars Surface Roving Vehicle,” Document 760-76, let Propulsion Laboratory, Pasadena, California; July 1972.*

6. V. F. Anthony, “Motion Control Requirements for Planetary Surface Roving Vehicles,” Document 760-78, Jet Propulsion Laboratory, Pasadena, California; July 1972.*

7. We thank Leonard Friedman for helpful comments on this point.

*JPL Internal Documents.

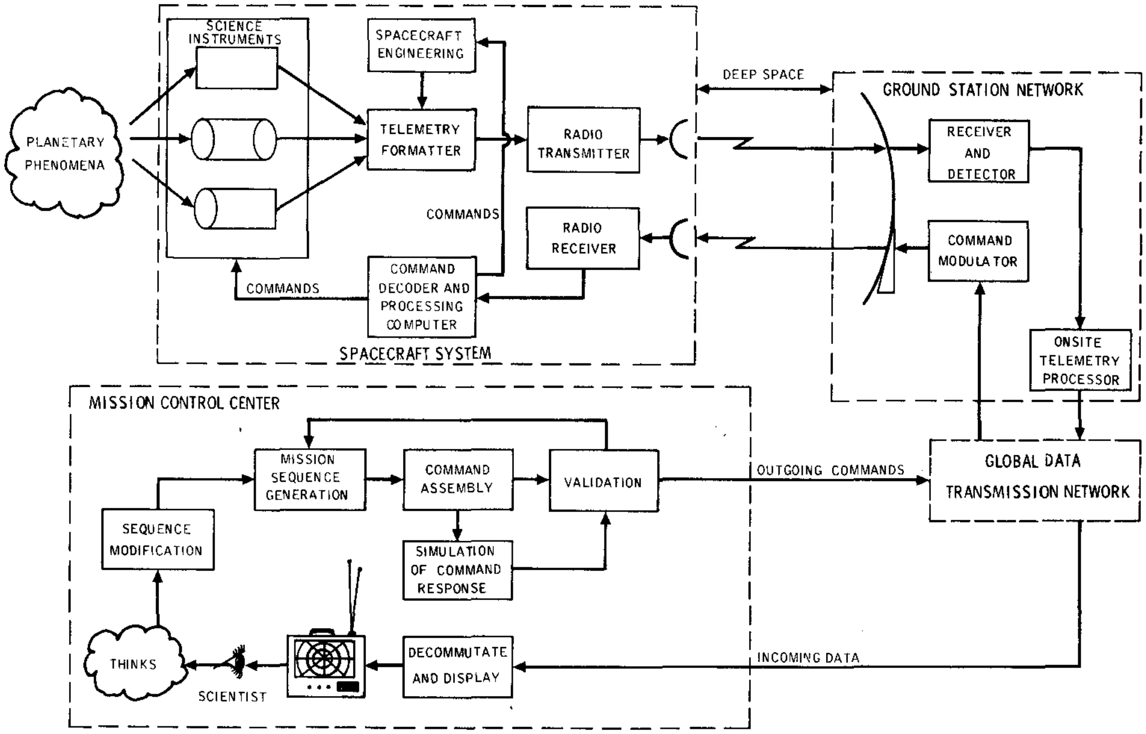

Figure 1. Spacecraft Control System

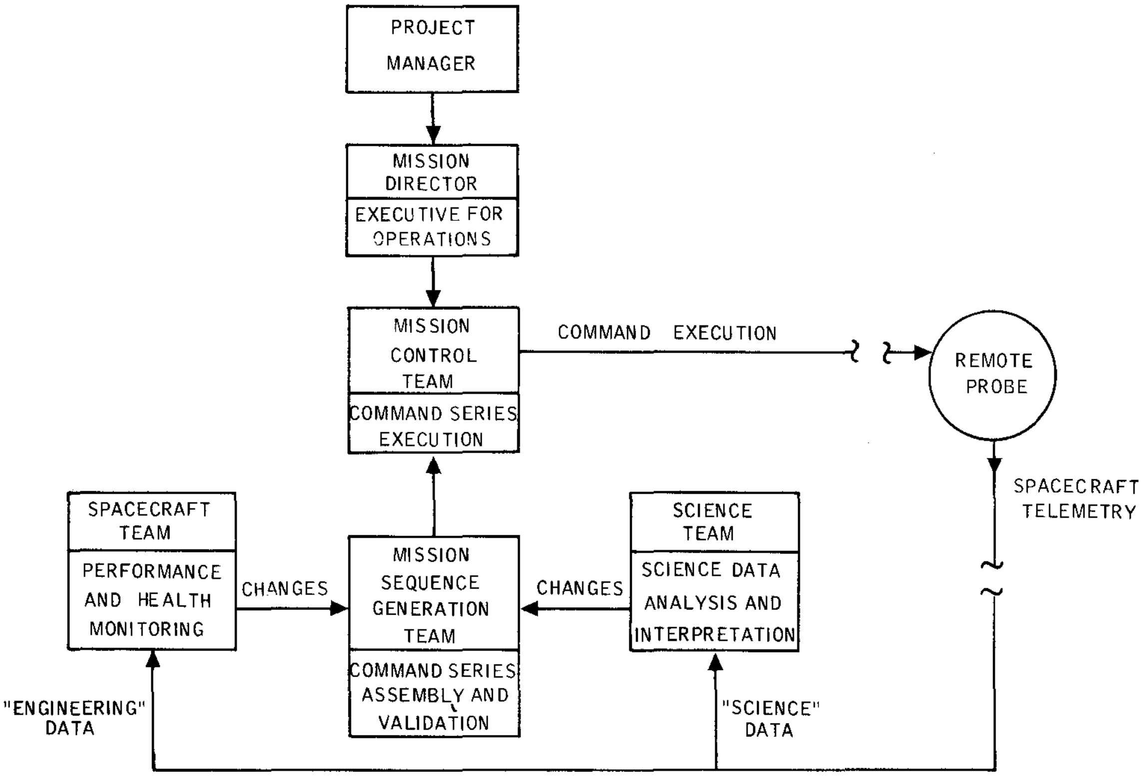

Figure 2. Mission Control organization and Functions