Anomaly Recovery and the Mars Exploration Rovers

J. Matijevic[1] and E. Dewell.[2]
Jet Propulsion Laboratory, California institute of Technology, Pasadena, California, 91109

The premise of the design of operations for the Mars Exploration Rovers (MER) is that the vehicles will drive each day. As a result, they will encounter some aspect of the terrain environment that cannot be anticipated or otherwise accommodated by the sequences linked onboard that day. The operations team then must correct the problem by planning then commanding the execution of a different drive the next day. Often other aspects of the operation on the surface of Mars: environmental changes, component degradation, errors in sequence design or execution, etc., lead to anomalies which must be addressed before normal operations can resume. The operational design that makes it possible to recover from a driving error each day also reduces the time needed to recover from anomalies. As an example of the efficiency achieved, less than 5% (about 30 sols out of 700 sols of operations) of the time on the surface has been devoted to recovery from anomalies for each vehicle. In this paper the major anomalies experienced by the MER rovers will be recounted and the streamlined approaches to addressing these problems described. The operational flexibility developed for these missions is also a function of the system design that anticipated a number of likely faults and conditions arising from uncertainty in sequence execution and environmental change. This design will be described as well as the considerations in operation that motivated this design. These considerations will likely be present in any future surface mission.

I. Introduction

As of this writing (MERA sol 780, MERB sol 760 where 1 sol = 1 Martian day = 24.6 hr.), the Mars Exploration Rovers (MER) have each spent over two years on the surface of Mars. They each have driven over six kilometers and encountered a variety of terrain conditions at the distinct landing sites of Gusev Crater (MERA, called Spirit) and Meridiani Planum (MERB, called Opportunity). They each have experienced a Mars year surviving and (sometimes) thriving in the change of seasons at the landing sites.

In the design of the rovers, variation in terrain, changing environment, and problems in operation from the surface of another planet were taken into account in the system. Although the implementation of the design is mainly 'single string' (e.g., one computer system, one transceiver, one inertial measurement unit, etc. per rover) a degree of functional redundancy in the design makes it possible to operate successfully with one failure.

Uncertainties in the Mars surface environment and past experiences with landed missions, lead the project team to treat a MER mission as a temporary resource for scientific investigation. Early projections of a prime mission of 3 months and perhaps (with luck) an extended mission of 3 months thereafter led to an operations team concept that treated any day on Mars as a precious asset. As such the plan for each day would contain combinations of science observations and engineering activities[1].

A long-range planning team in MER operations prepares outlines of future activities taking into account assumed progress from data acquisition and changes to location of the vehicles. This 'strategic' plan guides the operation through the creation of near term objectives for vehicle activity. Each day a separate team (on each vehicle) prepares a set of sequences of commands that implement the objectives for a short period (1-3 days commonly) of the strategic plan. This 'tactical' plan is a set of sequences prepared from 'scratch' based on the best knowledge of the state of the vehicle, with only the constraints of times of communication and resources of data storage, time and energy available. All sequences require development, assembly, and verification prior to uplink on the following

[1] Team Chief for Rover Engineering, Mars Exploration Rover Project, 4800 Oak Grove Dr., Pasadena, California, MS 264-422, Member of AIAA.
[2] Deputy Team Chief for Rover Engineering, Mars Exploration Rover Project, 4800 Oak Grove Dr., Pasadena, California, MS 264-422.

day. After that uplink, the process begins again with the report of the results of execution; that is, a new vehicle status is established and reported to the leads for developing the plan for the following day.

When a day's activity reveals that a problem has occurred, the tactical planning leads meet to outline the strategy for resolving the problem and continuing, as possible. At times the anomaly may be as simple as a sequence that did not fully execute. Often when that sequence was intended to move the vehicle, the on-board protection may have determined that no safe path could be found to continue a move toward the goal. The response in the next day's tactical plan would be to replan the intended motion and continue toward the goal. If the anomaly shows that a component has indicated a failure to complete a required operation or has not shown expected performance, additional data may be requested or (in support of same) an engineering test scheduled. Typically other parts of the rover are unaffected by the component anomaly, so the beginnings of corrective action (e.g., an engineering test) is simply added to other planned activities in the next plan. Sometimes little if any data were received from the execution of the last plan. This is in itself an anomaly and a recovery response is planned in the next plan. This recovery plan emphasizes getting data and generally little else. Typically, when this recovery plan works, the next plan resembles one of the prior described responses. Lastly, when the anomaly analysis suggests a persistent component problem, an anomaly team is formed and (perhaps) a multi-sol investigation is initiated. This takes the corrective action out of the realm of a tactical response and often requires engaging experts to help in the diagnosis and recovery.

The following describes the problems that the two vehicles have experienced and the process used to resolve these problems. This description is preceded by a brief overview of the system design with particular emphasis on the flight software system that enables the problem detection and recovery. An overview of the ground operation for the MER rovers is also provided.

II. System Design

The Mars Exploration Rover (MER) is a solar powered, six wheeled driven, four wheeled steered vehicle designed for operation on rugged terrain. Each vehicle carries an instrument complement consisting of a pair of narrow angled cameras (16deg FOV) each with independently actuated filter wheels (called Pancams), a miniature thermal emissions spectrometer (miniTES), an instrument deployment device (IDD) to carry the four instruments: an alpha particle x-ray spectrometer (APXS), a Moessbauer spectrometer (MB), a narrow field (8 mrad) microscopic imager (MI), and a rock abrasion tool (RAT). The vehicle has six additional cameras: wide angled (120 deg FOV) base body mounted (called Hazcams, paired front and rear on the vehicle) and intermediate angled (45 deg FOV) mast mounted (called Navcams). The mast holds both the Navcams and the Pancams while providing a periscope for the miniTES.

The instruments and vehicle equipment (motors, cameras, power bus, etc.) are wired to an electronics card cage called the rover equipment module (REM). The main computer is built around a RAD-6000 CPU (Rad6k), RAM and non-volatile memory. The non-volatile memory is implemented in a combination of FLASH and EEPROM.

Energy is collected from the solar array. This energy is channeled along the system power bus that is supported by two Li-ion batteries. Power not required to support loads is channeled to recharge the batteries. The batteries support loads drawing power in excess of that supplied by the solar panel. Power in excess of loads and recharge required by the batteries is channeled to an external shunt radiator. The regulation and distribution of power is managed by the battery control boards (BCBs), one for each battery and independently powered by the batteries. The batteries also supply power to the mission clock with an alarm clock feature, programmable by software.

Temperature sensitive components are located in the warm electronics box (WEB) which is an insulated compartment forming the body of a rover. Thermal control is primarily passive, with waste heat from electronics stored in the WEB during the day that radiates from the WEB during the night. Thermostatically controlled survival heaters and radioactive isotope heating units (RHUs) provide supplemental heating.

Communications are provided by two distinct systems: at X-band, a Small Deep Space Transponder (SDST), and two Solid State Power Amplifiers (SSPA), supported by a body-fixed, monopole Low Gain Antenna (LGA) and a High Gain Antenna (HGA) steered in azimuth and elevation; and at UHF, a transceiver, supported by a body-fixed, monopole antenna.

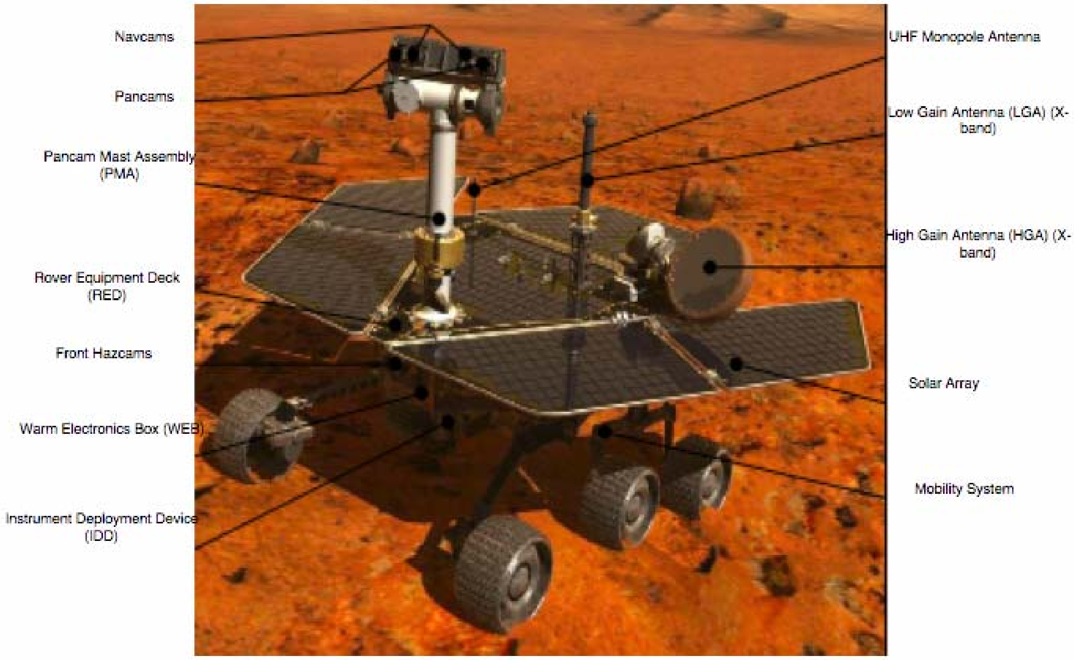

Figure 1. The Mars Exploration Rover

A. Flight Software Design[5]

The autonomous operation of the flight software maintains the vehicle in the state needed to receive and act upon commands, execute sequences of commands when available, and collect and format data for transmission. Separate software modules handle certain engineering functions of power and thermal monitoring, power on/off of components, conduct of communications, management of the alarm clock, memory management, control of process execution, device health status, and performance of sequence control. Payload functions are managed in other software modules which acquire images, process science data, control actuators, power instruments, and coordinate multiple actuations as necessary to drive the vehicle or deploy the IDD. All payload functions are executed under sequence control with both timed actions and event-controlled operations available.

Due to energy limitations expected at times during the mission, the flight software was designed to support wakeups (i.e., boot of the CPU) and shutdowns as part of normal operations. A wakeup is scheduled once each day when energy production from the solar array can support the load associated with the CPU and supporting electronics. This wakeup is called the solar array wakeup that occurs at the production of about 2A from the array that persists for at least 10min. The BCBs determine that the energy production from the solar array meets the 2A criteria and then can initiate a wakeup of the CPU and supporting electronics. A shutdown is controlled by parameters, established to ensure a power and thermal balance for the vehicle on any day when operations (including wakeups and shutdowns) are not otherwise commanded. A wakeup or shutdown may also be scheduled by sequence.

Communications are scheduled through the use of timed events maintained in an onboard communication windows table. At any given time this table contains about 6 weeks of timed events (windows) when uplink (commands transmitted to the vehicle) or downlink (telemetry transmitted by the vehicle) is planned to occur. The table contains both X-band and UHF communication events in a given period.

At each shutdown by flight software, a time is chosen for the next wakeup. This wakeup may be a sequenced wakeup, the beginning of a communication window or an autonomous event scheduled by the flight software. Choosing the nearest (in time) event, the flight software writes the time to the alarm clock, enables the clock to perform a wakeup, and then performs a shutdown (opens the power-on switch). At the next wakeup time, the alarm clock, independently powered by the batteries, can issue a wakeup through the BCB. The mechanism the BCB uses to wakeup the CPU and electronics is the same as the solar array wakeup: a power-on switch is held closed for a time period sufficient to allow the flight software to initialize and reinforce the power-on switch.

While in operation, the flight software records progress in execution through the generation of three types of telemetry: event reports (EVRs), engineering data (EH&A) and data products. When exceptions in processing occur, often all three types of telemetry play a role in announcing the occurrence and documenting the circumstances leading to the exception. The flight software has several levels of response to exceptions. At the lowest levels, an EVR (warning) may be written in the record to note the occurrence. At the next level a fault may be declared and the ongoing process or sequence ended. EVRs, EH&A, and perhaps a fault data product are generated for the record. Other processing will typically continue. At the highest level of response, the flight software will autonomously reboot when an unrecoverable problem is encountered. This is a fatal condition. Due to the seriousness of the problem, there is not time to document conditions at the time of the fatal. The flight software addresses this problem by temporarily storing a small number of EVRs in EEPROM during execution so that these can be recovered after the reboot has been accomplished after a fatal. The reboot puts the flight software into autonomous operation and processing continues as described above.

A fatal may also occur if a watchdog timer, used in monitoring progress in execution, is not refreshed. Once flight software has completed wakeup, there is one watchdog timer that must be refreshed by flight software 3 times in each interval of execution (RTI). Failure to perform this function leads to a reboot. Prior to wakeup, a watchdog timer tracked by the BCB must be reset within 4 minutes of power-on. If the flight software does not perform this refresh function the BCB will power-off the CPU and electronics.

One additional autonomous feature of the flight software was added during MER surface operations. This feature, termed 'deep sleep', causes flight software to remove the batteries from the power bus, either autonomously or by command. This operation mode can be invoked once the sun no longer illuminates the solar arrays. In this mode only the mission clock is powered; all other loads on the power bus are removed. This mode of operation ends when the sun illuminates the solar arrays. A time placed in the alarm clock function of the mission clock that is later than the time of illumination of the solar arrays is honored, causing a wakeup of the flight software.

III. Anomalies

The main anomalies on the MER vehicles are described in some detail in the sections that follow. For convenience, these are grouped into flight software anomalies, hardware anomalies, environment induced anomalies and an anomaly of uncertain origin. All anomalies were observed through the execution of the functions of the flight software but later attributed to purely a software fault, hardware failure or an unaccounted change to the environment. The final category is comprised of the reset anomaly on Opportunity, an anomaly that has not been fully explained.

Conspicuous by it absence is any discussion of the FLASH anomaly on the Spirit vehicle. This remains the worst problem experienced by either vehicle. The two weeks for recovery given a problem which impacted the normal operation of the flight software and thereby communications remains a testament to the capabilities and dedication of the MER operations team[3].

B. Software Anomalies

1. The race condition, initialization counter

The first instance of this problem appeared on Spirit sol 131. At the time of the downlink of telemetry on the UHF-band, the vehicle was unexpectedly in autonomous operation; that is, no sequences were active on board. Further a fatal exception was noted in the EVR log in the telemetry. In particular, the initialization module for flight software generated the EVR. After some investigation of the software code, the problem was seen to be a vulnerability that occurs when the initialization module was attempting to increment the counter of the number of times of initialization. This counter resided in non-volatile memory. Writing to this memory required permission from a separate software service that managed access to the memory. In the instance on sol 131, between the request and the grant of access to write to the memory location, another software module had requested, been granted access and had written to non-volatile memory. The initialization module, finding it could not write the initialization counter, declared an exception that resulted in a fatal condition.

Within the structure of the flight software on the vehicles, all processes time-share the use of the single CPU. There is no guarantee that the three actions desired by the initialization module (i.e., request, being granted write access, and writing to memory) would occur contiguously. In this case, the vulnerability to the fatal would be viewed as a function of the number of processes in operation at the time of the write of the initialization counter. Clearly one corrective action could be to restrict the action of other software modules during this period of vulnerability. In part for some payload modules, that action was implemented. But the vulnerability was viewed as so limited in time (a few microseconds to perform the three actions desired by the initialization module albeit within about a 4 minute window during initialization), the anomaly team reviewing the problem decided on issuing only an advisory for this vulnerability. The planning teams, noting the vulnerability of the motion of the IDD to an unexpected fatal exception, restricted use of the IDD from the 4-minute window during initialization. Otherwise, the likelihood of recurrence was deemed so slight as to not warrant further action, including consideration of any flight software change to attempt to correct the problem. This is understood as a race condition between software modules: a race that the initialization module won for many initializations (over 560 at that time) and many sols since this occurrence.

Recurrence proved not so unlikely as the problem occurred again on Spirit on sol 209 and then on Opportunity on sol 596 and sol 622. At the last recurrence, the fatal occurred during initialization in preparation for an afternoon UHF communication window. The flight software response caused the loss of that communication window. As a consequence, on sol 622 the downlink team received no telemetry, leaving the recovery team to sift through many possibilities for the problem. Eventually, at the next uplink window on sol 623, commands were issued and accepted by the flight software, and sequence control was reestablished. Sol 624 was a sol of normal operations.

Due to the delay in recovery of a sol after the sol 622 event, the anomaly team recommended enforcing a 'keep out zone' for science operations after wakeup. Due to energy considerations however, the recommendation was only enforced during the wakeup prior to an afternoon UHF pass. This ensured that a recurrence would not jeopardize the return of engineering and science data needed to plan for the next sol.

2. Another race condition, imaging interface

This anomaly occurred on Spirit on sol 136. There were no data returned on the sol 136 afternoon UHF pass. Plans for that sol and the next sol were such that there was a morning uplink communication window and afternoon UHF passes with two relay assets on sol 137: MGS (Mars Global Surveyor) and then from ODY (Mars Odyssey Spacecraft) on sol 137. In an attempt to receive data at the earliest opportunity on sol 137, a command was issued to Spirit to make the uplink communication window a 2-way uplink/downlink communication session on X-band. This action failed to result in a return of data.

When data were received from both passes, an analysis of the data showed that the vehicle was unexpectedly in autonomous operation and a fatal exception was noted in the EVR logs. The fatal exception occurred when imaging sequences were deactivated prior to the preparation time for the afternoon UHF communication window (imaging cannot occur simultaneously with UHF communications due to EMI). The imaging had not completed when the sequence deactivation commands were issued. The imaging software module, performing the image read/data write operations, was starting an image read when a command (part of the deactivation) was issued to power off the hardware resource used in image acquisition. The software module, seeing the return of an error when attempting to access the hardware resource, declared an exception that led to the fatal exception response by the flight software. Since the fatal exception occurred during preparations for the afternoon UHF pass, the communications window could not be executed by the flight software (the window time was in the past when the reboot after the fatal exception occurred).

Further, when the reboot due to the fatal exception occurred on sol 136, the status of the system at the time of the fatal exception was not completely saved in memory. In this case, the position of the HGA antenna was not saved.

This played a key role in the attempt on sol 137 to conduct communication through X-band. An X-band fault occurred since flight software had no knowledge of the position of the HGA when attempting to move the HGA in position for communications. This was the explanation for the failure to receive data during the attempt at X-band uplink/downlink communications.

A modification to the deactivate sequence process was recommended after this incident. A directive to the imaging module to end execution was added prior to the power off of the image acquisition hardware in the deactivation process. Once implemented, there has been no further occurrence of this problem. Normal operations were resumed on sol 138.

3. Corrupted command: conjunction test

As a test of the degradation in command receipt experienced on the surface of Mars during solar conjunction (defined as the period at which the Sun-Earth-Mars angle is less than 2deg), commands were issued to the two vehicles each day during the period from MERA sols 244-255 and MERB sols 224-234. These commands were NO_OP commands, dummy commands issued to enable acknowledgment of receipt in the EVR logs generated by flight software. As the test progressed, more commands (at each command session tens of NO_OP commands were issued in groups) were corrupted upon receipt as noted in the EVR logs. Finally, on MERB sol 229 a corrupt command caused a fatal exception. In analysis of the EVR logs and subsequent review of the command upload software module, a command with sufficient number of errors could be executed and cause a stack clearance to occur. This would explain the fatal exception.

Although the problem in the software could be corrected, the recommendation was to simply suspend commanding during the remainder of the conjunction period on both vehicles. Normal operations were resumed on sol 235 (first day after conjunction).

4. Exception in evaluation of a DDI during mobility

The design for sequence control supported by the flight software includes the definition of a defined data item (DDI) that can serve as a global flag for control of sequence execution. However, flight software clears this global flag after it has evaluated the flag in a conditional statement. There is only one such global flag, so the definition and then evaluation should take place as a contiguous process. In the example of a set of drive sequences on sol 449, a DDI was defined in two sequences that were executed in parallel. The execution of this set of parallel sequences eventually resulted in an evaluation in one sequence following the definition in another.

The fatal exception that occurred was reported in the afternoon downlink in the UHF pass on sol 449. The drive that was implemented in the sequences of sol 449 ended prematurely and none of the imaging typically performed after a drive was acquired (the reboot after a fatal exception places the vehicle in autonomous operation, so no imaging sequences (as an example) commanded with the drive will be executed). This fact plus the common practice at this time of the mission to avoid working weekends resulted in additional sols for recovery. Normal operations did not resume until sol 453.

A recommendation after this event was to restrict the use of a DDI to one sequence at a time. A flight software modification is planned to change the strategy of evaluation of a DDI, thereby removing the vulnerability to a recurrence.

5. Upload fault during forward link commanding

There had been no plan in the MER mission to use the forward command link capability of the UHF communication system: commands would be issued on X-band and telemetry would be returned on a combination of the X-band and UHF systems. The longevity of the MER missions and the exigency imposed by periods of absence of X-band coverage led the operations team to develop this capability while operating the two MER vehicles.

The first demonstrations of the facility for commanding a MER vehicle through ODY relay were conducted during the prime mission (within the first 3 months) and periodically within a year after landing. These demonstrations were implemented through the issuance of real-time commands that verified functionality of the command link. The first attempt at sending a full sequence load through the relay link was on MERA sol 603. This was a single sol plan with 15 sequences. All were received successfully. The second attempt on MERA sol 605 was a 3-sol plan with 41 sequences. Again all were received successfully, although in the EVR log there were a number of warning messages associated with processor loading. As understood at the time, these warning messages were a benign consequence of the CPU failing to 'keep up' with the combination of forward link commanding conducted simultaneously with downlink through the UHF transceiver.

These demonstrations helped to prepare the operations team for a formal readiness test conducted in flight on MERA with commands issued through the forward link system and executed on board. The first forward link upload was a 3-sol plan containing 53 sequences commanded on sol 640. These were received successfully but again with warning messages as seen on the prior sequence loads. The second forward link upload was a 2-sol plan with 43 sequences commanded on sol 644. After receiving 12 of the 43 sequences, there was a fatal exception in the flight software. On analysis of the EVR logs and a review of a simulation of the events on the upload of sol 644 in a test bed for the MER vehicles, a combination of the number of sequences and a dropout in the forward link due to geometry between ODY and the vehicle resulted in the fatal exception. The remainder of the readiness test was cancelled and recovery to normal operations was achieved on sol 646.

This incident was followed by an effort to develop a workaround, robust to the number of sequence and the possible geometry-induced variability in communication link performance between the ODY relay and the surface vehicle. After testing with a MER test bed, a strategy for forward link commanding was developed: more 'padding' (null characters) was introduced between sequences in the command files transmitted by ODY and the number of sequences transmitted was reduced to not more than 25 with allowance for a few additional real-time commands. The increased padding in the command files reduced the processor contention during command decoding, verification and storage while telemetry was being processed and transmitted. The reduced number of sequences was consistent with the protocol overhead and the command rate possible with the forward link. This strategy was demonstrated successfully by test sequence loads issued on sols 747 and sol 755. In both instances, 25 sequences plus an additional real-time command were issued. Forward link commanding was successful in practice when the uplinks on sols 773, 775 and 777 during a period of no X-band coverage for Spirit resulted in operations on sol 774 through sol 779.

C. Hardware Anomalies

6. Stuck-on Heater

Upon receipt of Opportunity’s first overnight UHF pass on sol 2 at 03:30 LST, the power team reported that the nighttime loads from sol 1 22:54 LST to sol 2 3:30 LST were ~0.5Amps larger than predicted. The subsequent pass showed that the additional load had remained on until ~10 LST, dissipating ~176 W-hrs. An anomaly team convened to develop a possible explanation and recommend a resolution. The highest likelihood fault was that a Rover Power Distribution Unit (RPDU) load was unexpectedly powered on. Due to the size and the on/off times for the load, the IDD heater #1 was identified as the source of the problem. The on/off times for the load corresponded to the predicted thermostat box switch times and to the temperature rise recorded by the temperature sensor on the Microscopic Imager (MI), located close to IDD heater #1 when the IDD is in a stowed configuration. Due to the design of the thermostat box, the IDD heater was not powered on during the day (~10LST to ~23:00LST). This prevented overheating in the event that an IDD heater circuit was stuck on. This design also suggested a mitigation strategy for the energy loss due to the stuck-on heater circuit; remove the heater from the circuit at night.

The rover was still completely functional as designed. However, the energy drain from the heater would reduce the energy available for science activities. As winter approached, when less energy was available, the activities on the spacecraft would be much more limited, and the survival of the spacecraft would be at issue. The solution (removing the heater from the circuit at night) was implemented by arranging to remove the batteries from the power bus. By removing the batteries from the power bus, the Battery Control Board (BCB) was also not powered, leaving only the mission and alarm clocks powered. At dawn, the BCB are powered (awakened) simply by having sufficient light impinge on the solar arrays (about 0.2 Amps of current generated by the array). The net savings would be ~180 W-hr/sol. The team implemented this solution, termed ‘deep sleep,’ for the first time on sol 101-102. ‘Deep sleep’ was permanently enabled on sol 206, so that it was the default state unless temporarily disabled for the night. One drawback was that survival heaters for the miniTES and REM, which are normally left on during the night, were taken off-line. ‘Deep sleep’ is currently used on average every other night on Opportunity.

The use of 'deep sleep' to mitigate the 'stuck-on' heater energy drain enabled an extended mission for Opportunity. It came at a price however. With batteries removed from the power bus, no survival heater for the miniTES could be powered over a night of ‘deep sleep’. The colder temperatures on the miniTES (routinely below the acceptable flight temperature limits) due to 'deep sleep' have undoubtedly contributed to a degradation of this instrument on Opportunity (see the section on resets below). This degradation has not been experienced by the miniTES on Spirit. Due to the generally colder temperatures at the Gusev landing site and the absence of a significant energy drain, Spirit has never used 'deep sleep'.

7. RF drive actuator

Spirit's drive from Bonneville crater to the Columbia Hills was planned to satisfy the two-fold objective of reaching potentially a new geological target for investigation and enabling the vehicle to spend the winter at northerly tilts (of benefit for vehicle safety for a solar powered system). The drive was nearly 3 km and needed to be completed in a period from sol 86 through sol 156. The drive was accomplished by driving on 50 of the possible 70 sols. This was a remarkable accomplishment for Spirit at that time in the mission, but it was achieved at the cost of increased current draw seen on the right front drive actuator. The drive actuator for the MER rovers was designed as a geared, lightly lubricated system. Each drive actuator required on average approximately 0.4A during motion of the vehicle. The current draw can spike when the vehicle is engaging an obstacle but these spikes are generally less than 1A for a fraction of a second. By the end of the period of the drive to Columbia Hills the right front drive actuator required nearly 1A while the vehicle was in motion with spikes over 1.2A. Further the current draw for this actuator had increased exponentially during the last 10 drives. At that rate, further degradation would lead to a drive actuator that could not be supplied sufficient energy to move. It was estimated that this loss of the drive actuator could occur within the next 100m of travel.

An anomaly team was formed to review the data and the possible options for recovery. The team proposed stopping the drive at the base of the hills and performing a series of motor diagnostics. A relatively flat area was chosen for these diagnostics that consisted of small forward and backward motions with the actuators warmed prior to the operation. The only slight improvement in performance of the right front actuator after these tests resulted in the team developing a strategy for conserving further use of the actuator. The strategy involved driving the vehicle 'backward' while dragging the right front wheel. Periodically, the right front wheel would be used to help correct the error in direction created by having the vehicle move while dragging the right front wheel. With this strategy, Spirit climbed into the hills and continued its mission. Over the period of the succeeding 200 sols, the combination of reduced use of the right front drive actuator, driving backwards, warming the drive system before driving and simply standing still, usually at a spot where in situ investigations were conducted, resulted in the right front drive actuator eventually returning to operation (i.e., current draw) comparable to that seen on the other drive actuators. This remains the case at this writing. The problem with the actuator was attributed to the flow of the lubricant in the gearbox. The persistent driving in one direction 'starved' the gearbox of lubricant causing the higher, anomalous current draw. Warming the actuator, reversing the direction of driving and not driving every few days was sufficient to correct the problem.

In this period of about 200 sols after the problem was detected only 4 sols were devoted to engineering tests designed to diagnose the problem and practice the strategy of driving while dragging a wheel. Science observations were conducted and the vehicle was driven (albeit for short distances) in a nominal fashion while the anomaly team was performing analyses and developing the strategies to conserve right front actuator usage.

8. RF Steering actuator

On Opportunity on sol 433 in the middle of a drive, the right front steering actuator stalled. In the period from sol 358 to sol 433, Opportunity had traveled nearly 3 km, driving on about half of the days in that period. The stall occurred without warning; that is, the steering actuator was used nominally on the prior segment of the drive. On the drive segment in question, the steering actuators were being positioned for a turn when the steering actuator current peaked to the preset safety limit (2A).

An anomaly team was quickly formed to review the images acquired during that day and the data logs at the time of the incident. A test was proposed and executed on the next sol. The test involved commanding the steering actuator clockwise then counterclockwise a few degrees at several voltage settings to test for any motion. There was little movement and that motion only at the highest voltage setting during the test. The anomaly team concluded that there was an obstruction in the actuator, likely at the first planetary stage. As such, the motor alone has little torque to move the obstruction. The team recommended leaving the actuator at its current position and operating nominally. Fortunately, when the motion of the actuator was stalled, the right front wheel had been steered only about 7 deg from the nominal straight-ahead position. In subsequent drive operations, the planning team pivots about the right front wheel using larger arc turns rather than turns-in-place to position the vehicle.

Normal operations of the payload resumed on sol 434, and driving continued within a few days (sol 437). At this writing, the strategy of pivot turns around the right front wheel and variants such as 'K' (or three point) turns continues to be used.

9. IDD azimuth actuator

On Opportunity on sol 654, the instrument deployment device (IDD) was commanded to unstow from its location below the WEB shelf. The first motions of the IDD involve the movement of the elevation and elbow actuators that uncouple a hook on the elbow assembly from a support roller. After this operation is completed, the shoulder azimuth actuator moves the hook away from the roller in preparation for a subsequent movement of the elevation and elbow actuators to move the IDD away from all support structure below the WEB shelf. The shoulder azimuth actuator did not move in this sequence causing the unstow operation to fail.

An anomaly team was formed to consider tests that could be conducted to both move the shoulder azimuth actuator and collect data during that motion. Changes in control parameters that define the movement profile of that actuator were proposed and planned for execution on sol 659. There was little motion on that attempt. Changes to these same control parameters, effectively allowing for more current to be supplied for longer periods during the motion, were attempted on sol 660 and sol 661. There was little motion on either attempt but what motion was seen suggested that the motor resistance had increased significantly from the nominal value calculated during prior (and successful) unstow operations. A motion with a control parameter of a higher resistance was planned for execution on sol 666. This motion succeeded and the detailed data verified that the motor resistance had doubled. With the increased resistance value as a control parameter, the IDD was unstowed successfully on sol 671. Subsequently, a science investigation using the IDD and planned at the location of Opportunity on sol 654 was carried out. The motion of the IDD and, in particular, the performance of the shoulder azimuth actuator was somewhat restricted and a little unpredictable. However, the experience in operation at this site gave the operations team a performance characterization that was used at other sites in the vicinity. This experience also provided the foundation for the operation of the IDD from this point forward in the mission.

A past characterization test conducted with the type of motor used in the implementation of the shoulder azimuth actuator revealed a failure mode in which the motor resistance had doubled due to an open wire in one of the windings. There are two windings in this brushed motor. A second open wire would cause the motor to fail. The likely cause of degradation in the motor winding is thermal cycling. This actuator is on the 'stuck-on' heating circuit and the temperature transient is nearly twice the scale of that of all but one other actuator on the IDD and elsewhere on the vehicle as a whole. The other actuator is the elevation actuator of the IDD that shares the problem of the 'stuck-on' heating circuit.

The anomaly team discussed how the vehicle could drive with the IDD with a degraded shoulder actuator. If the IDD was stowed under the WEB shelf and the shoulder azimuth actuator failed, the IDD and all instruments for in situ measurement could no longer be used in the mission. If the IDD was unstowed, how much driving could be allowed and what terrain could be traversed so that the vehicle motion did not damage the IDD or (at a minimum) cause a loss in calibration for the IDD. A position, with the wrist and turret suspended over the solar array (the 'hover stow' position), was analyzed and shown to tolerate drive excursions that involve differential motions of the vehicle of about 4 cm. Several targets were accessed by Opportunity driving with the IDD in the 'hover stow' position. For longer drives, an evaluation of the terrain to be traversed cannot ensure that the differential motion constraint is satisfied (e.g., not all areas planned to be traversed in a long drive can be imaged). Instead, the anomaly team agreed that the IDD must be stowed in the usual position below the WEB shelf for any long drive. Since the open wire anomaly in the winding of the shoulder azimuth joint is understood to result from thermal cycles, it was recommended that the IDD be stowed prior to the drive then unstowed immediately after the drive. The deepest thermal cycle would then be avoided with the IDD stowed below the WEB shelf. The first successful drive using the stow/drive/unstow strategy occurred on sol 731. At this writing, there is additional analysis planned for the eventuality when the IDD can no longer be stowed below the WEB shelf (i.e., another open in the winding for the shoulder actuator occurs causing loss of that actuator; degradation or loss of the elevation actuator). Driving may be problematic at best at that time.

Throughout the period from sol 654 to sol 731, each sol contained nominal operation of other payload elements. And after sol 671 IDD usage was scheduled.

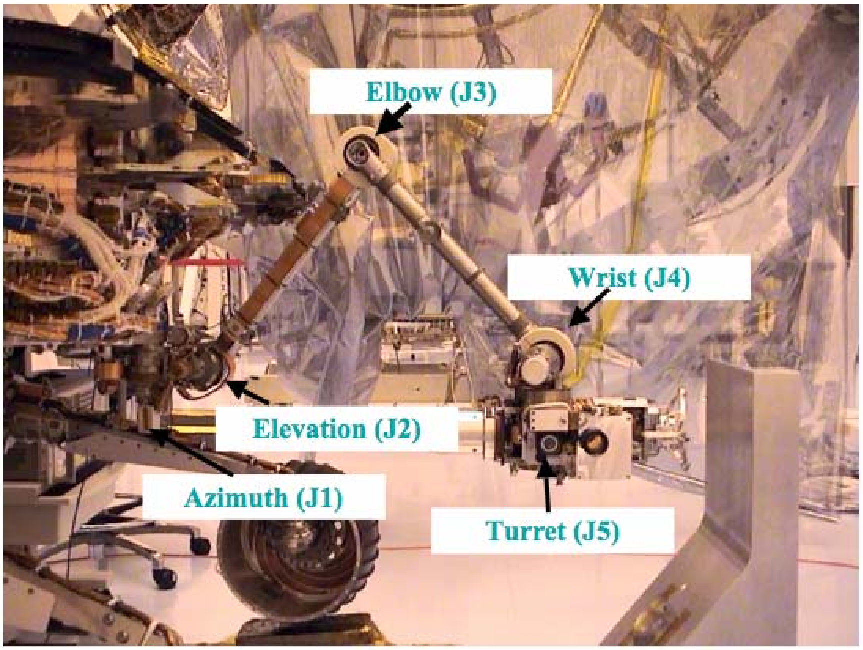

Figure 2. Instrument Deployment Device (IDD)

D. Environmentally induced anomalies

10. Clock Fault

On the morning of sol 628, Opportunity encountered a dust storm that significantly increased the tau value (optical opacity of the sky) from 0.82 to 1.8. This increased 'haziness' in the sky, caused the sunlight to be more diffuse, and thus provided less power to the solar arrays. The result was that both the BCB wakeup and solar array wakeup occurred later than the uplink team expected. The plan on the evening of sol 627 was to invoke 'deep sleep', and schedule a sequenced wakeup at 07:40 to re-enable the miniTES heaters. The previous 2 sols' BCB wakeups were at 07:23. However, due to the dust storm, the BCB wakeup on sol 628 occurred just after 07:40. Thus, the 07:40 alarm clock occurred while the batteries were still off-line, and nothing responded to the wakeup request. The next time the rover had a wakeup was at the solar array wakeup at 10:44 LST, at which time the alarm clock time was in the past, and so the flight software deactivated all sequences and thereafter was in autonomous operation. The next time given to the alarm clock was for the start of the afternoon UHF pass, which was the time at which the operations team learned a problem existed. After the UHF pass, the alarm clock set the rover to wakeup at 18:30LST to do a 'wakeup for deep sleep' followed by a shutdown bringing the batteries off-line (i.e., 'deep sleep' is the default in autonomous operation). However, the on-board communications table had an AM pass onboard, as the plan for the evening of sol 628 and morning of sol 629 was to disable deep sleep for a single night activity. This resulted in the alarm clock being set for that AM pass and, as a consequence, a clock fault occurred again that following morning of sol 629. This fault did not affect the flight software that was already in autonomous operation from the fault on sol 628. The only impact from this second fault was that due to the residual effects of the dust storm a late solar array wakeup at 11:18LST caused the loss of an X-band communication window that was supposed to occur from 10:30-11:00LST. A somewhat mystified operations team later worked out why no data were received from that scheduled X-band window. Since this was part of a 3 sol weekend plan, and the vehicle was otherwise in a power positive (and eventually) understood state, the operations team waited until the following Monday plan (sol 630) to resume nominal operations.

The uplink team decided to not wake up after deep sleep to enable the miniTES heaters for the next few sols until the dust storm had dissipated to prevent missing the BCB wake up time after deep sleep. Once the skies cleared, the planning team established a guideline to schedule wakeups 45 min to 1 hr after the BCB's come online.

11. 'Potato' Rock

On sol 339, Spirit's rover drive planners had planned a drive up Husband Hill in route to an area deemed scientifically interesting and an ideal place to take panoramas of the hill traversed thus far. The progress had been slow, with the rover encountering high slip and surrounded by a rocky hillside. At the first turn command in the drive sequence, the rover was commanded to turn in place 15 degrees in 17.5 seconds. In the last 2 seconds of the turn, the right rear wheel current spiked to 1.8 amps (nominally the wheel actuator draws about 0.35 amps) and the drive motor actuator stalled. A rock was lodged between the inner ring of the wheel and the actuators as could be seen in the imagery transmitted on the sol 339 afternoon UHF pass. On the following sol, the operations team dislodged the 'potato' rock (so called due to its size and shape) by spinning the right rear wheel in the opposite direction of the sol 339 drive. The rock was dislodged from the actuator, but remained inside the wheel. The next challenge was to get the rock out of the wheel, and the operations team decided to try dragging the wheel in an arc. On sol 343, the right rear wheel was straightened, and then the drive and steering actuators were disabled (a temporary conditional setting in flight software). Two small 0.3 meter arcing drives were attempted using the remaining 5 wheels. However, the rock remained in the wheel. On sol 344, the rover was commanded to back down the hill (all actuators enabled this time), but the rock remained in the wheel. On sol 345, the rover was commanded to turn in place, to drive back down the hill, and to perform a final turn in place. The sequence of motions completed successfully, but the 'potato' rock was still in the wheel. Finally, on sol 346, two drives and one more turn in place were commanded. On the afternoon UHF pass, images confirmed that the 'potato' rock was out of the wheel. Nominal operations resumed for the weekend plan on sols 348-350.

During development of the rover mobility system, the design considered the potential of small rocks kicked up by vehicle motion into portions of the drive mechanism. Rock jamming in the mobility mechanisms had occurred on prior rover systems with the result that in the implemented MER system there was additional external clearances around the wheels and the drive actuator mechanisms were located within the wheel well. The wheel wells were not 'closed out' with structure due to weight considerations for the vehicle as a whole. Despite these provisions, the geometry of the vehicle and the loose rocks and regolith of the terrain on Husband Hill made this rock jamming possible.

Figure 3. ‘Potato’ Rock Stuck in Left Rear Wheel

12. Embedding in terrain

On sol 446 Opportunity drove into a dune. The drive was planned for 90m with almost all of the drive executed with the vehicle driving backwards (bogie first). This drive was the 36th segment of a drive across the Meridiani plain that had already covered over 3km in distance south from the vicinity of the landing site. As were many of the prior drives, the drive on sol 446 was conducted 'blind' with few safety checks employed during execution. The environment had enabled this type of driving since a generally flat, featureless terrain was presented to the operations team on each of the past drives, including the terrain imaged on sol 439 and used in planning for the sol 446 drive. Also ripples in the regolith and small dune features seen in previous terrain images had posed no hazard for the mobility system for Opportunity. The ripples and small dunes were unremarkable in the images of the terrain used for planning the sol 446 drive and these features were not considered any hazard for the vehicle. About 40m of the drive on sol 446 had completed when the embedding in the dune began. The vehicle was slipping and embedding itself in the dune while completing the final 50m of the drive. A final turn at the end of the planned drive to achieve the best position for return of data in a communication session was not completed and only at that time was a fault declared by the flight software. The material transition was over 2 meters and the rover climbed upon a dune over 30cm in height. The final images returned on the UHF pass on sol 446, showed the wheels about 70 percent buried.

Upon review of the data from this drive, the operations team was mainly concerned that the last segments of the drive and the turn had resulted in little change in the rover position. Also, on inspection, the wheel cleats were 'caked' with regolith. During development of the MER mobility system, tests showed that the cleats played a significant role in the engagement of terrain, modifying the ground and creating mechanical 'purchase' that allowed forward motion of the vehicle to occur. Wheels without cleats experienced slip in otherwise benign, sandy terrains. Would Opportunity be able to escape this dune without exposed cleats and regolith piled to the top of the wheels?

An anomaly team was formed with individuals who had participated in the test and eventual qualification of the MER mobility system. Added to this team were science team advisers with understanding and test experience with vehicles driving in earth-based terrains. Tests with test bed vehicles driving out from beach sand were promising: the test bed vehicle drove easily out of the terrain beginning with wheels buried completely in the sand. From soil tests conducted during the earlier phases of the mission, the team surmised that the Meridiani regolith was likely a less cohesive mixture of material than beach sand. A variety of materials were mixed with the objective of making something that behaved like Meridiani regolith. Such a mixture of materials could be used as a simulated regolith in test with the test bed vehicle. Eventually, a mixture of diatomaceous earth, mason clay and sand proved to have a consistency like that seen in soil tests on Mars. This material had the added benefits of being generally available and certified for personnel exposure in test facilities. An area in the MER test facility was prepared with this mixture and the test bed vehicle was driven into it. A series of tests were conducted with the wheels buried at the levels seen on Mars with Opportunity. In each case the test bed rover drove out of the material, although with difficulty. About 50m of wheel motions were required to move the test bed vehicle about 2m out of the material with most of the motion occurring at the end of the drive.

With these tests completed, the operations team began to drive Opportunity out of the dune. Driving the vehicle out of the dune would initially be in short movements, designed to determine if there was any chance that vehicle motion would cause further embedding into the dune. The first moves (a wheel straightening and a few meters of driving) began on sol 461 and sol 463. Subsequent drives through sol 468 resulted in less than 10cm of forward motion per each drive of about 8m or roughly 99% slip. After sol 469, 8m then 12m then 20m of motion was commanded until the vehicle was extract from the dune. In total, 2m of forward travel required almost 200m of commanded motion with the last meter of travel occurring on the final sol.

In this period, the operations team developed then tested strategies that provide protection from embedding events on future drives. These strategies included: measured current draw versus forward motion, use of visual odometry to determine slip in forward motion (applied on any drive greater than 5m), incremental turns to avoid large displacements of regolith by the wheels, analysis of terrain to predict successful traverse and drive length reduced to 30m or less (i.e., driving within the imaging range of the vehicle). The dune of this embedding was 35cm tall: the largest dune formation seen in the terrain prior to sol 446. Imagery from this location and thereafter showed that dunes of this size were increasingly common as the vehicle moved south. Reconnaissance imaging would be a part of any drive planning from this point forward.

While driving was limited during this period from sol 446 to sol 484, science observation continued throughout this time. The mobility analysts and rover drive planners on the operations team developed and carried out the testing program in the test bed. They developed then implemented the drive strategies that have resulted in about 1.5km of successful driving since that time.

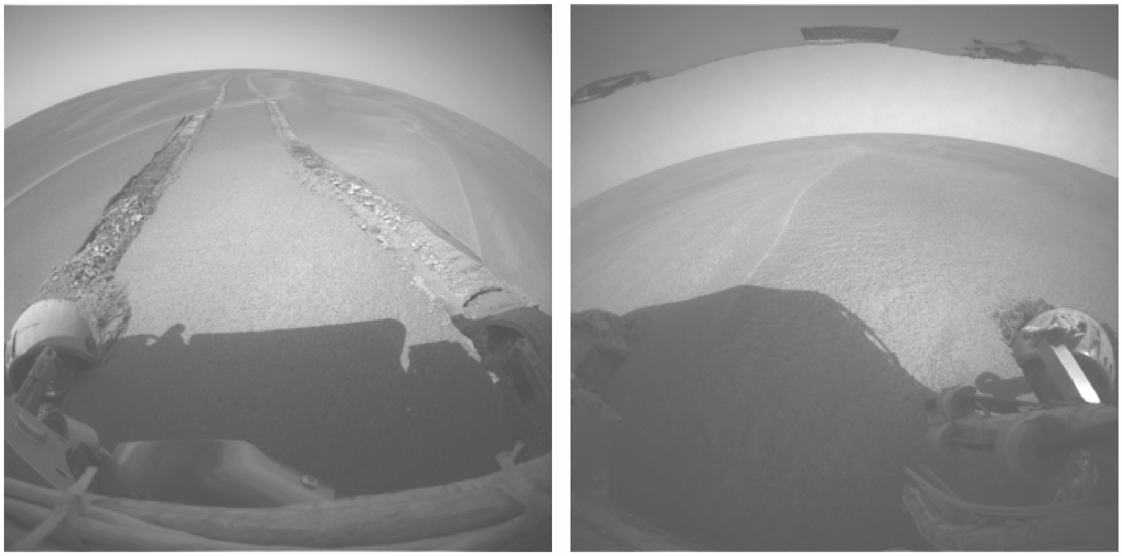

Figure 4. Front and Rear views of Opportunity embedded in a dune

E. Anomaly of Unknown Origin

13. Resets

On each of three sols (sol 440, 563 and 610) Opportunity experienced an unusual reset event: an unplanned shutdown then wakeup. Each of these events occurred during the execution of a miniTES observation and little data were reported at the time of wakeup from each event. In the case of the sol 440 reset event, the system response was consistent with the response to a flight software fatal exception; that is, an autonomous reboot occurred. However, in this event, there was no EVR reported for the fatal exception and those EVRs stored prior to the exception (as a short history of execution) shed no light on the cause. The cause of the reset was reported in an EVR received on the subsequent wakeup as the watchdog timer failing to be refreshed. This is the watchdog timer used as a self-test while flight software is executing. In the cases of the sol 563 and sol 610 reset events, the system response was consistent with a watchdog timer prior to wakeup not being refreshed and a reset was initiated by the BCB. In both cases, a miniTES observation was initiated (a sequence was activated) prior to completion of the initialization process by the flight software. Such sequence activation is allowed by the flight software and has been successfully executed many times during the mission of both rovers.

At the time of the sol 440 reset event, an anomaly team was formed to consider the possible causes. Without data related to the cause of the event (why was the self-test watchdog timer not refreshed?), the focus of the investigation was the consideration of the possible collateral effects associated with the continued use of the miniTES. Through test on the vehicle, the PMA and associated motor controller were determined to not be causes of the problem (other payload items successfully used these components and the associated electronic interfaces subsequent to the sol 440 reset). The electrical interface to the miniTES instrument was reviewed for possible failure effects. This review showed that no short or open in the wiring to the instrument would cause a severe cascade, damaging other equipment on the electronic board providing the interface to the miniTES instrument or other portions of the REM. Other possible causes principally associated with flight software in operation (i.e., a VME bus error, a spurious uplink command, a software controlled reset, an under-voltage in a power converter unit, and a BCB watchdog timer induced reset) were all reviewed and discounted due to a lack of evidence of occurrence. A self-test watchdog timer failure could be the result of a software module 'hanging' in execution. However, the software modules in execution at the time were routine health and maintenance tasks of flight software, motor control and the miniTES payload interface module. All these modules had performed many similar operations prior to this point (and thence thereafter) without incident. A code review of the miniTES payload interface module revealed no particular vulnerability for 'hanging' in execution.

After a recovery sol where component states affected by the reset event and interfaces to the PMA were tested and shown to be working nominally, normal operations (without observations by the miniTES) were resumed on sol 442. After the review of both hardware and software was completed, the anomaly team concluded the miniTES was not a significant risk either to the instrument itself or to other rover equipment. On sol 487 routine miniTES observations were again scheduled on Opportunity.

At the time of the sol 563 reset event, the anomaly team considered the reset more serious than the sol 440 reset event. The BCB initiated reset could be associated with a problem in the rover power system or in the BCB itself. Further, since a BCB initiated reset does not automatically reboot the flight software system, the next software wakeup could be an alarm clock timer wakeup scheduled 26 hours in the future. The wakeup process loads 26 hours from the time of last wakeup into the alarm clock as a precautionary value intended to allow the rover system to recharge batteries in autonomous operation prior to attempting to wakeup and initiate flight software. This design covers a variety of possible power system anomalies, all of which could result in a flight software failure to wakeup successfully as indicted by a BCB initiated reset. If this value in the alarm clock is not changed by a subsequent successful flight software wakeup, the rover system could spend a sol or more not communicative with the MER operations team and not under sequence control. This is a serious problem that in the past cascaded into the loss of a Mars surface mission.

Fortunately, in the sol 563 reset event (as was also the case in the similar sol 610 incident) the automatic solar array wakeup was the next wakeup event. The flight software in autonomous operation after a successful wakeup honored subsequent communication windows including the afternoon UHF window on sol 563. At the time data were reported from sol 563, the system had returned to predictable operation. On the succeeding sols a slow recovery to nominal operation was planned. Execution of sequences that exercised the PMA, the mobility system and various health and maintenance functions of the flight software were performed. No problems were observed with any of these functions and nominal operations (without the miniTES) resumed on sol 571.

The anomaly team had few leads for determining a source of the problem. There were no data recorded during the event or reported from the time of the event. The timing was such that the miniTES observation was only initiated before the BCB induced reset occurred. The team reviewed the software implementation of the BCB but could determine no particular problem with the watchdog timer function. The rover system performed as designed but the question remained: what caused the reset. As recommendations, the anomaly team requested that if the miniTES was used again, its sequences should be executed outside of the period of wakeup initialization of the flight software. This operation restriction may cause a subsequent reset to occur in the manner of the sol 440 reset event. Also, the team recommended that an attempt be made to replicate the circumstances of the system response to the reset on one of the MER test bed systems. There was no miniTES instrument built for use in a ground test facility. A simulator was provided during MER software development that included an interface to the RS422 standard used on the electronic board that provided a power and data connection to the miniTES instrument in the flight configuration. The miniTES payload interface module had been developed using this simulator and there was no record of a problem such as this software 'hanging' during execution: a possible reason for a watchdog timer to not be refreshed. Perhaps applying an external stimulus at this interface to the simulator in the test bed could cause a reset response like that on Opportunity.

In lieu of the results from any test bed investigation, the miniTES was still considered no significant risk either to the instrument itself or to other rover equipment. On sol 570 a miniTES observation was again scheduled on Opportunity. With restrictions (e.g., scheduling miniTES sequence execution after completion of wakeup initialization of flight software), miniTES observations were regularly scheduled after sol 577. The operations team assembled a standard recovery plan on the chance that a reset could occur in the future. This plan reduced the time of recovery to a single recovery sol that contained the tests necessary to determine that nominal operations could resume. Due to successful execution of the miniTES from sol 577 through sol 598, all miniTES restrictions were lifted thereafter.

The decision to allow unrestricted miniTES usage proved to be wrong with the recurrence of the sol 610 reset event. As was the case on sol 563, the BCB reset the flight software prior to wakeup due to a watchdog timer failing to be refreshed. A miniTES sequence was just activated when the reset occurred. Due to the similarity with the sol 563 reset event, there was no reconstituting of the anomaly team. In contrast to the response to the sol 563 reset event, the MER operation team applied the streamlined recovery plan on sol 611 and Opportunity was ready tor resume nominal operations on sol 612. No further miniTES observations were scheduled until sol 656. At that time only one observation was allowed per week based on an agreement of the principal investigator and the project manager for MER operations. This restriction was lifted on sol 728 when the test bed investigation was completed.

While the operations team was conducting the recovery leading to normal operations on Opportunity, an engineer, consulting with members of the anomaly team, prepared one of the test beds with break-out-box and signal generator at the interface between the miniTES electronic simulator and the REM data and power system[4]. Tests of interruption in the power and data signals to the electronic simulator were performed while the flight software was in operation consistent with the use of the miniTES during the reset events. After many trials, a test case resulted in the triggering of a 'hanging' of the flight software due to the interruption of signal while the flight software was in the midst of the initialization process. The 'hanging' caused a BCB initiated reset, similar to the sol 563 and 610 events. At other times (generally after most of the initialization had been completed) the interruption of signal resulted in the declaration of a fatal exception. This fatal exception was accompanied by a warning EVR that pointed to a portion of the processing in the miniTES payload interface module. This processing used the 'zero path difference' (ZPD: a mean estimate of the size of the spectra generated during an observation by the miniTES) as a parameter in the calculation of compression of the data produced by a miniTES measurement. In the test case the ZPD value was zero causing a software cascade leading to the 'hang'. Repeating the test with versions of this software module with a fixed (though artificial) nonzero ZPD value or without execution of the portion of the processing involving the ZPD in data product production resulted in no occurrence of the software 'hang' and the subsequent BCB initiated reset or fatal exception.

So why should zero ZPD values ever occur? Why should this occur on Opportunity and not Spirit? Why should this occur at this late stage of the mission? The explanation for this may be traced to an incident and subsequent corrective action that occurred only on Opportunity and only after the miniTES experienced a pronounced degradation. At that time (after sol 394 on Opportunity) the instrument routinely produced short interferograms (a representation of the spectral output of the instrument) that the flight software rejected for production into data products and transmission to earth. The analysis of the problem at that time (from a consensus of a combined contractor, university and MER engineering team) was that the instrument had suffered degradation in the servo-mechanism used in control of the scan mirror within the instrument. The cause was likely the result of repeated cold cycles associated with the use of the operation of 'deep sleep'. This technique of 'deep sleep', required to continue the mission of Opportunity beyond the prime mission period, exposed the miniTES to thermal cycle below the allowable flight temperature of -40degC. The resolution at the time of the sol 394 and subsequent incidents was a software change that accepted short interferograms in the production of miniTES data products. Such short interferograms can have zero ZPD value in an anomalous case.

Unfortunately, the data collected from the reset incidents on Opportunity do not include a data product from the miniTES at the time of any reset event. Associating the test case with the flight incidents may be an explanation for the resets but cannot be viewed as root cause for the incidents. The test case did reveal a vulnerability that was corrected with an operational restriction on the type of data product produced by a miniTES observation. This restriction was imposed after sol 728. In addition, the restriction that miniTES sequences should be executed outside of the period of wakeup initialization of the flight software was enforced after sol 728. There have been no subsequent incidents.

IV. Conclusions

The MER operation process was developed to allow the rovers to drive every sol. Since at least the terrain environment changes when the vehicle moves, the results of a drive require evaluation before the rover can be safely commanded to drive again. This requirement led to scheduling command interactions and telemetry return every sol during the mission. This process (daily evaluation and planning on the tactical timeline) proved to be beneficial in also training the operations team to respond on a tactical shift to the uncertainties in command execution and to changes to the assumptions about continued rover operation. This is the training necessary for recognizing and responding quickly to anomalies.

The major anomalies experienced by the rovers are described and the response of both the tactical operations team and, when necessary, an anomaly team are detailed. In general, the operations team continued payload operations within a few days after every one of these events. The anomaly team, cognizant of the need to continue the science observations of the missions, proposed engineering tests and changes in vehicle procedures that fit with nominal science operations.

V. Acknowledgements

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration

Glossary of Acronyms:

WEB Warm Electronics Box
IDD Instrument Deployment Device
PMA Pancam Mast Assembly
REM Rover Equipment Module
BCB Battery Controller Board
MER Mars Exploration Rover
EVR Event report
ODY Mars Odyssey Spacecraft
MGS Mars Global Surveyor Spacecraft
HGA High Gain Antenna Assembly
LGA Low Gain Antenna Assembly
DDI Defined Data Item
RPDU Rover Power Distribution Unit
miniTES Miniature Thermal Emissions Spectrometer
SDST Small Deep Space Transponder
SSPA Solid State Power Amplifier
RHU Radioactive Isotope Heating Unit

Keywords:

Rovers, sol, fault, autonomous operation

VI. References

Papers and Reports

[1]A. Mishkin, D. Limonadi, S. Laubach and D. Bass, "Working the Martian Night Shift, The MER Surface Operation Process", to appear in IEEE Robotics Automation Magazine Special Issue on Mars Exploration Rovers, June 2006

[2]T. Neilson, "Mars Exploration Rovers Surface Fault Protection", Proceedings of the 2005 IEEE Conference on Systems, Man, and Cybernetics, Waikoloa, HI, Oct. 10-12, 2005

[3]G. Reeves, T. Neilson, "The Mars Rover Spirit FLASH Anomaly", 0-7803-8870-4/05 2005 IEEE paper #1425, Proceedings of the 2005 IEEE Conference on Systems, Man, and Cybernetics, Waikoloa, HI, Oct. 10-12, 2005

[4]K. Fleming, “Mars Exploration Rover Mini-TES/PSAB Software Interface Fault Testing”, JPL internal report, Feb. 7, 2006

[5]G. Reeves, J. Snyder, “An Overview of the Mars Exploration Rovers’ Flight Software”, Proceedings of the 2005 IEEE Conference on Systems, Man, and Cybernetics, Waikoloa, HI, Oct. 10-12, 2005