Corresponding Author: Mishkin

Tracking Number: 58994

From Prime to Extended Mission: Evolution of the MER Tactical Uplink Process

Andrew Mishkin* and Sharon Laubach[ε]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

To support a 90-day surface mission for two robotic rovers, the Mars Exploration Rover mission designed and implemented an intensive tactical operations process, enabling daily commanding of each rover. Using a combination of new processes, custom software tools, a Mars-time staffing schedule, and seven-day-a-week operations, the MER team was able to compress the traditional weeks-long command-turnaround for a deep space robotic mission to about 18 hours. However, the pace of this process was never intended to be continued indefinitely. Even before the end of the three-month prime mission, MER operations began evolving towards greater sustainability. A combination of continued software tool development, increasing team experience, and availability of reusable sequences first reduced the mean process duration to approximately 11 hours. The number of workshifts required to perform the process dropped, and the team returned to a modified "Earth-time" schedule. Additional process and tool adaptation eventually provided the option of planning multiple Martian days of activity within a single workshift, making 5-day-a-week operations possible. The vast majority of the science team returned to their home institutions, continuing to participate fully in the tactical operations process remotely. MER has continued to operate for over two Earth-years as many of its key personnel have moved on to other projects, the operations team and budget have shrunk, and the rovers have begun to exhibit symptoms of aging.

Nomenclature

CoUGAR	=	Command and Uplink Generation and Review
MER	=	Mars Exploration Rover
MGS	=	Mars Global Surveyor
RSVP	=	Rover Sequencing and Visualization Program
SAP	=	Science Activity Planner
SOWG	=	Science Operations Working Group

I. Introduction

THE Mars Exploration Rover (MER) mission successfully landed two mobile geological laboratories on the Red Planet in early 2004. As originally envisioned, MER was intended to operate the two robotic rovers—named "Spirit" and "Opportunity"—on the surface of Mars for approximately 90 Martian days (or "sols") each. The

* Principal Engineer, Planning and Execution Systems Section, MS 301-250D
[ε] MER Integrated Sequence Team Chief, Planning & Execution Systems Section, MS 264-422

rovers' lifetime was assumed to be limited due to expected buildup of dust on the solar arrays, which in turn would reduce the energy available for rover science activities and traverse. This constraint resulted in the design of an intense tactical operations process capable of achieving the defined mission objectives within the presumed short mission duration. Due to the inherent non-determinism in the execution of rover traverse and instrument placement in the poorly modeled planetary surface environment, and distinct from most robotic deep space missions, MER surface operations needed to be reactive, with the plan for each sol responding to the results achieved on the prior sol. For the MER prime mission, these requirements led to a seven-day-a-week multi-shift 18-hour command-turnaround process slaved to the Mars local time at each landing site, which planned one sol's worth of rover activities during each planning cycle. The two operations teams, one for each rover, literally lived on Mars-time, with each team in a different Martian time zone.

As the MER surface mission unfolded, it became clear that the rovers were likely to survive and remain functional for far longer than originally anticipated. The operations process in place would require modification to enable sustained operations over an indefinite period, while reducing overall operations team size and lowering operations costs. Process and software tool re-design eventually reduced the duration of the tactical operations process to 8 to 10 hours, which has enabled a single-shift planning cycle and consolidation of several team positions for the current extended mission. We describe the evolution of the tactical process to its current form, as well as the logistical challenges of working on a 5-day-a-week Earth-time schedule, including planning up to 3 rover sols of activity during one operations team work shift. For example, due to Earth-time/Mars-time phasing, uplink planning for the next sol must sometimes proceed before the current sol's activities have completed. In such cases, the team must plan in "restricted" mode, such that rover driving or instrument arm placement is permitted only every other sol, to provide the opportunity for the success of rover motion activities to be assessed prior to planning the next motion. As of this writing, the MER extended mission continues, with both rovers conducting science observations on a daily basis, operated by a team of scientists and engineers one-third the size of the operations team present during prime mission.

II. Drivers and Attributes of the Original MER Operations Design

A. Limited lifetime

During the development of the MER project, the lifetimes of the two rovers were expected to be short. The "warranty" (i.e., the mission lifetime requirement) for the rovers was specified as 90 sols, or 90 Martian days. Assuming successful landings, energy models indicated that the rovers would survive as long as 140 sols or so. Data gathered during the Mars Pathfinder mission's 83 sols in 1997 had shown a dust buildup on the solar arrays of about 0.2% per sol. A combination of this dust settling and the changing Martian seasons (particularly for Spirit, located further from the equator) would in time leave the rovers with too little energy for traverse, and eventually not even enough for overnight survival heating. Some forecasts indicated that the effect of the dust would be so pronounced that the rovers would be capable of achieving more in the first 30 sols of the surface mission than in the remaining 60 sols. This lifetime constraint put a premium on every command opportunity.

B. Reactive operations

For most robotic deep space missions, a single command load may be prepared weeks or months in advance of use, and govern the actions of the target spacecraft for a like period of time. For a planetary surface mission in which rover traverses and otherwise interacts with the terrain, there may be more relative uncertainty in the vehicle's position after an attempted drive of 3 meters than in a free-flying spacecraft's position after a million kilometers of travel through space. A rover may slip while driving across the surface, or hit a buried rock and make no further progress. The Spirit and Opportunity rovers are capable of autonomous hazard detection and avoidance; they may reach the desired target location, but due to the need to avoid a hazard unseen by ground operators, may take longer to reach that location. And the choice of what target to approach next, and the determination of where it is safe to drive or deploy the instrument arm, depend on the images from the latest rover position. All of these factors force the operations team into a reactive mode: to plan what a rover should do next, the results of what the rover just attempted must be known.

C. Limited communications opportunities and downlink bandwidth

The solar-powered rovers would not have the energy to transmit continuously through their X-band direct-to-Earth links continuously, and would also be limited by heating of the transmitter over time. Although UHF relay through the MGS and Mars Odyssey orbiters was incorporated into the mission design, these links had yet to be proven, and were presumed to have long latencies. Therefore, the mission developed the concept of "critical" data, defined as that data required to support the next planning cycle. Critical data must fit into the X-band session, while most science data not needed for operations planning would be relegated to the UHF relay links. The capacity of the X-band link would be about 20Mbits per sol.

Additionally, two factors limit the number of opportunities to transmit commands to the rovers: again, lack of energy to keep the high gain X-band receiver powered continuously, and the need to negotiate use of the Deep Space Network with all other deep space projects, including other spacecraft at Mars (such as the other rover).

D. Time Delay

With roundtrip communications time delays between Earth and Mars of 6 to 40 minutes, depending on relative orbital position, direct teleoperation or "joysticking" of the rovers would not be feasible. This constraint drives the use of stored command sequences for controlling the rovers, and the requirement for the rovers to be capable of autonomous surface navigation and hazard avoidance.

E. Every-sol commanding

Given the presumed limited surface lifetime and non-determinism of rover activities, the key to achieving the MER mission objectives was to command the rovers as often as reliably possible. For nominal surface operations (after rover deployment from the lander), the selected frequency of commanding was once per Martian sol. In this approach, the rover would downlink its critical data in mid-afternoon each sol, after completing all major traverse or robotic arm activities. The ground would then uplink the next command load the following Martian morning. Since the rover would be largely dormant overnight (except for in-situ instrument integration and UHF relay sessions), the rover state as downlinked in the afternoon would be a good predictor of the state at the time the new command load would be received onboard. Even with every-sol commanding, over 1% of the entire surface mission would be consumed per command cycle.

F. Mars-time

How does one schedule operations teams to operate spacecraft slaved to another planet's diurnal cycle? The mean length of a Mars day—at about 24 hours and 40 minutes—is frustratingly close to an Earth day. MER chose to put its operations teams on the same schedule as the rovers: Tactical team workshifts started 40 minutes later each day. By kicking off the operations process as soon as the Mars afternoon downlink was received on the ground, the team always had the maximum possible number of workhours available between the downlink and the next uplink. (Early in the mission, virtually all of these workhours were consumed.) Working Mars-time also ensured that shift handovers occurred at the same point in the planning process, avoiding the necessity of cross-training the entire team in multiple roles.

Of course, since the two rovers were to land on opposite sides of the planet, the operations teams for each rover would be living in different Martian time zones.

III. Operations process for the MER prime mission

At the start of the surface mission, the MER operations process was described as the "overnight" tactical timeline, since the process was performed primarily during the Martian night while the rover slept. The process executed according to a strict schedule (see Figure 1), to ensure that the command load would be ready for uplinking before the communications opportunity arrived. The key steps of the tactical process are described below.

A. Receipt of downlink

After completing its final traverse or instrument placement for the sol, the rover transmits its critical data during a communications session up to about an hour in duration.

B. Engineering downlink assessment

As the data is received on the ground, engineering team members immediately begin assessing the health of the rover, determining that no anomaly has occurred and that there are no contraindications to performing a normal planning cycle. Automated tools begin constructing data products from the downlink stream, including images. Calibrated images are due for delivery to the planning process within 30 minutes of the end of the downlink.

C. Science downlink assessment and science activity planning

Even before downlink begins, subgroups within the science team begin planning observations and partial activity plans contingent on the successful execution of the current sol's plan. When the downlink arrives, science team members assess the health of the onboard instrument suite, and the veracity of the instrument data products. The team continues activity planning, representing and targeting the activities in the SAP (Science Activity Planner) software tool. Depending on the received downlink, activities may be refined, modified, or thrown out and replaced. At 1800 Mars local time, the Science Operations Working Group (SOWG) meeting is convened. At the start of this meeting, a representative of the engineering team presents the status of the rover, and re-confirms the resources available for the next sol. Subsequently, proposed activities from the several subgroups are merged, rejected, and re-worked. Under the direction of the SOWG Chair, the team prioritizes activities for inclusion in the science plan, and for level of criticality to the next planning cycle. As it is being worked, the preliminary plan is repeatedly modeled for resource usage, including duration, energy, and data volume generated.

Members of the Integrated Sequence Team (IST) attend the SOWG meeting, and assess the feasibility of the preliminary science activity plan from two perspectives: 1) will the rover be capable of executing the plan on Mars, and 2) will the operations team have time to transform the proposed plan into a complete, validated command load in time for the Mars morning uplink opportunity. Among other evaluators, the Rover Planner engineer assesses the feasibility and safety of any traverses or instrument placement required to reach science targets requested by the science team.

The SOWG meeting must deliver the file defining the preliminary science activity plan for the sol by 2000 Mars local time. The bulk of the science team is released from the tactical process at this point.

D. Activity plan refinement and validation

The center of activity now shifts to the sequencing room, where members of the IST, instrument sequencing experts from the science team, and the SOWG Chair continue developing the rover command load for the sol. The team integrates engineering activities into the preliminary plan, formally schedules all activities, and models the plan in detail for resource consumption. Activities that do not fit within the available resources will be deleted from the plan. The SOWG chair represents the overall science team for any such intra-science decisions, guided by the established activity priorities; this approach allows for on-the-spot resolution of conflicts to keep the time-constrained process on track. The activity planner must produce a conflict-free plan within the available rover resources in time for the scheduled review meeting. Higher priority activities are scheduled first, so that if time runs out, only the lowest priority items will have been lost.

While the overall activity plan is in work, the Rover Planner uses the Rover Sequencing and Visualization Program (RSVP) to design and simulate all rover traverse and robotic arm activities. RSVP fully represents the 3-dimensional terrain model derived from the stereo imagery returned from the rover. Due to the fidelity required to plan rover and arm motion activities, this part of the planning process is equivalent to developing a draft version of the rover motion command sequences.

E. Activity plan review

At 2345 Mars local time, the Activity Plan Approval Meeting (APAM) is convened. The team reviews both the overall activity plan and the rover motion plan. All documents and visual simulations of rover motion are projected from appropriate team member workstations, so that the entire co-located team can review them together. The team confirms that resources are within allocations, that activity-level flight rules have been satisfied, and that despite any changes the plan remains consistent with the original intent of the science team. Only very minor changes to the activity plan are permitted at the APAM, such as insertion of opportunistic remote science observations that will not impact onboard resources.

The APAM also serves as a shift handover for many team members. Team members coming on-shift carefully study the final activity plan, and ask questions during the review. After the meeting, incoming and outgoing team members in corresponding roles may confer one-on-one to ensure that the new team member has the situational awareness to take over.

F. Command sequence generation

As soon as the SOWG meeting ends, the instrument sequencing experts begin designing the command sequences to implement all instrument activities. Using RSVP, they work in parallel with the activity validation and rover motion planning tasks, to ensure that the sequences will be ready in time. However, until the activity plan has been approved, there is the risk that a given activity will be deleted from the plan, and therefore that any sequence built to execute that activity may never be used. This risk of wasted effort is small compared to the time saved by not waiting for the final activity plan to be delivered. Thus, sequencing begins well before the APAM, and continues for two hours after the APAM.

Following the APAM, a team member produces the "master" and "submaster" command sequences, which together embody the activity plan in the form of commands. These sequences command rover wakeups and shutdowns, and determine when each of the individual instrument activity sequences will kickoff.

G. Sequence integration and validation

At 0200, the team reviews the master and submaster sequences in detail. Any errors identified are corrected in real-time, and the modifications reviewed again by the team. All sequence developers must deliver their individual validated sequence files by the conclusion of this review.

The Sequence Integration Engineers now take on the task of integrating the 30 or more individual sequences into a single command load, produce the command files and review products, and perform both manual and automated cross-sequence flight rule checks on the full sequence set, while other team members conduct a software simulation.

H. Command review

The final review gate is the Command Approval Meeting, at which the team confirms that all validation and flight rule checks have been successfully performed, without discovery of any errors that will impact execution of the command load.

I. Transmission of commands to the spacecraft

At the Command Approval Meeting, the command forms are approved, together with the instructions specifying the exact order of transmission of the command files and associated real-time commands. These signed documents (together with the activity plan approval form these are the only paper documents required by the tactical process), are hand-delivered to the workstation from which the commands will be sent to the proper DSN station and transmitted to the spacecraft. Due to the time delays inherent in Earth-Mars communications, no immediate acknowledgment of command receipt is available from the rover. Instead, The DSN listens for a carrier signal "beep" tone at a specified time determined by commands contained in the new command load, thereby confirming that the new load is running onboard. If the beep is instead received at a second specified time, this will mean that the prior sol's command load is still running, and the new load was not properly received.

Figure 1. MER TACTICAL TIMELINE

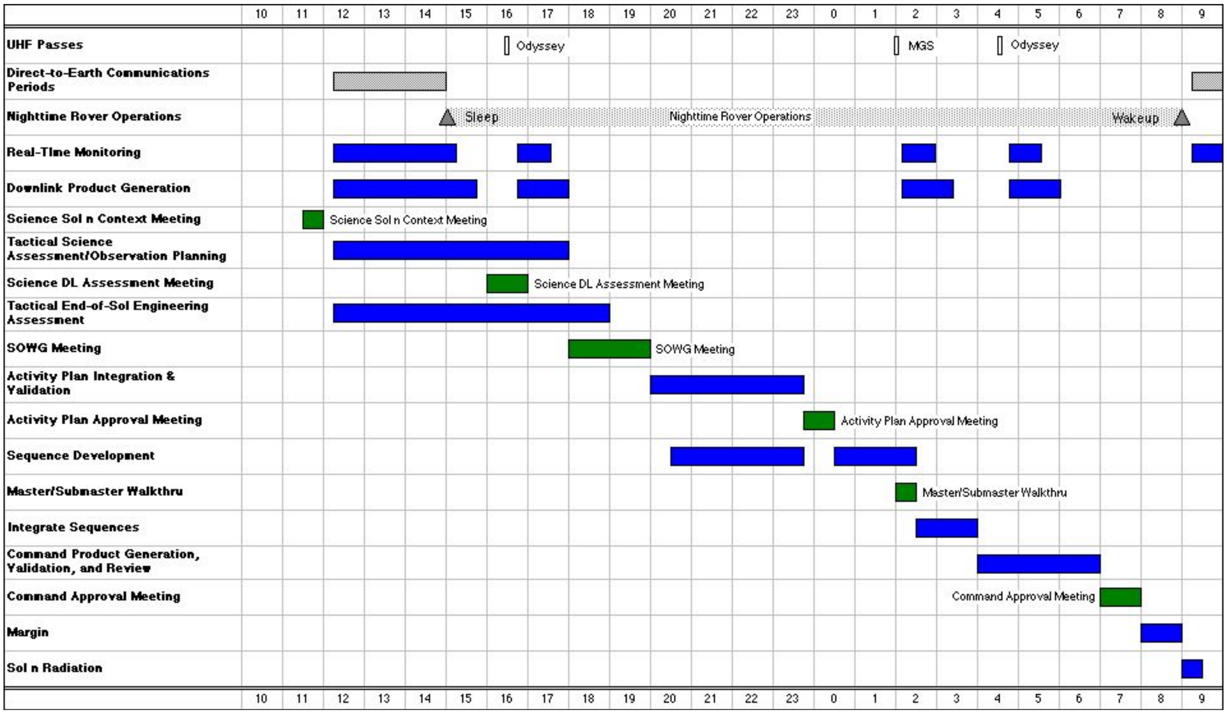

IV. Extended Mission 1: Returning to Earth

Living on Mars-time was no mean feat. It meant that the team would emerge from their shift at the same Mars local time each sol, but the corresponding Earth local time would rotate through the day—so if at the beginning of the prime mission someone's shift began in the early evening, a week later they'd be starting work in the middle of the night, and a week after that, mid-morning. Outside the MER tactical environment, the team had to deal with family, friends, and other tasks which still ran on normal Earth-time schedules, as well as attempt to sleep before they had to return for their next shift. On top of living on Mars-time, the shifts themselves were long and intense. Fatigue, despite countermeasures taken to prevent it, was always a potential issue. Additionally, team members felt increasingly isolated from the world outside MER's tactical environment. Thus, it was always clear that if the rovers continued to survive, the original operations design would have to change, in order to be sustainable.

The key change required for sustainability was getting the team off of Mars-time. Since the main driver for working Mars-time was the number of hours required for turning around a command load, the length of the tactical shift had to be reduced. Automation proved to be the enabling technology, supplemented by the accumulation of reusable instrument command sequences, and streamlining of team member functions as the team learned exactly what was needed. Together, these improvements shortened the tactical process just enough so that margin could be spent to return the team to a modified Earth-time schedule.

Due to Earth-time/Mars-time phasing, and to the fact that the rovers are still a wasting resource, the team could not simply change to a fixed Earth-time shift schedule. Instead, a scheme was devised which would maximize the commandability of the rovers, given the new tactical process duration, but still limit how much a team member's shift could move relative to Earth local time. Since the rovers' uplink and downlink times are still slaved to Mars local time, those milestones would still rotate through the Earth day. When the downlink time hit early enough in the Earth day, and uplink was late enough, the team could work normal hours. As downlink shifted later and later, the team's work day would have to start later, until the end of the tactical process hit the latest possible time the scheme would allow. At that point, the tactical process would begin before the downlinked telemetry had arrived and been analyzed. Since planning robotic motions required knowledge of the state of completion of prior commanded robotic motion (including terrain imagery), these types of motions could only be allowed every other sol. Hence, this period in the modified Earth-time schedule, when the tactical process began before the prior sol's telemetry had been analyzed, is termed "restricted sols". When the downlink and uplink times again rotate far enough through the Earth day, the team returns to normal, unrestricted planning. This cycle reoccurs every 37 Earth days.

The team continued to shorten the tactical process, primarily by augmenting their automation tools. One package which had a drastic effect on process length was the Sequence Integration Engineers' tools, called "CoUGAR tools" (for Command and Uplink Generation and Review), in order to automate the generation of the "master/submaster" sequences, integration of the science and engineering sequences, and generation of the uplink products and the review products as much as possible. Robustness checks were added, to help reduce the amount of time spent on review. In addition, a method was developed so that the Sequence Integration Engineers could "piggyback" on robotic simulation results already created by other team members, avoiding the time and risk of re-running those simulations. Another change was the decision to limit other kinds of modeling done at the command level, instead relying on resource modeling at the activity level to validate the sol's plan. This decision required that no significant changes be made to the plan downstream of the Activity Plan Approval meeting, but further reduced the length of the tactical process. Overall, these improvements, along with increased reuse of sequences and a better understanding of what activities would fit into a given sol's plan, shortened the tactical process to roughly 8 hours, and allowed some roles to be combined which now fit into a single shift. The new process duration was re-factored into the modified Earth-time schedule, allowing the schedule to continue to approach normal working hours over the seven-day workweek.

V. Extended Mission 1: Working 5 Days a Week

As it became clear that the rovers were continuing to survive far longer than the 90-sol prime mission, it was also burdensome to the team to continue to work weekends. The decision was made to allow the team to plan three sols' worth of activities on Fridays, and uplink the resulting 3-sol command load. Again, since the results of robotic motions over the weekend could not be known during planning, the weekends were treated as "restricted sols". The CoUGAR tools, designed for every-sol commanding, were significantly revamped to allow robust multi-sol planning without severely impacting the length of the tactical process. In the end, each additional sol in the plan added only an hour to the tactical process, including time for selection of activities for the additional sols, as well as review of the multi-sol products at various points in the process. A skeleton staff reviews downlinked telemetry over the weekend, and a small on-call team is designated to respond to anomalies as needed.

VI. Extended Mission 2: Distributed Operations

For the second MER extended mission, it no longer was feasible to keep the entire team collocated at JPL. The original operations design's near-independence from paper documentation facilitated the move to distributed operations, but it was unclear at first how well the collaborative tactical environment, so dependent upon interaction between team members, would translate. Initially, some science team members moved back to their home institutions and attended the SOWG meeting remotely. Given the success of this first move, in September 2004, the SOWG Chair experimentally participated entirely remotely, leading the SOWG meeting from another room and being available during the remainder of the tactical process only over the teleconference line. The process was slightly changed to accommodate the remote SOWG Chair; for example, a webcam was set up to broadcast the information projected onto the screen in the sequencing room during review and approval meetings, and selected intermediate products were moved to a remotely-accessible area. Again, the experiment was a success, and now the SOWG Chair is nearly always at their home institution. Finally, using the infrastructure prepared for the remote SOWG Chair, the remaining team members from outside JPL were able to move back home and deliver their sequences or other products remotely.

VII. Extended Mission 3: Sustaining Operations and Turnover

Many of MER's operations team members had been involved in the development of the MER rovers, and had expert knowledge of how the rovers' systems functioned. As MER moved into its third extended mission, a year after the prime mission ended, these experts began moving on to other projects to return to spacecraft development. MER's tactical process shifts, though shortened, could still be intense, and the robustness of the plans and short duration of the tactical process, as well as appropriate response to anomalies, was partially dependent on expert operators.

In addition to the pull of other projects on MER's experts, the realities of budget reductions forced a reduction in the size of the operations team, further straining the operations process. As personnel rolled off the team, "corporate knowledge" about past operations strategies was sometimes lost. An unforeseen consequence of the very successful automation effort was that new personnel sometimes did not learn the details of the process “under the hood”, so when unusual circumstances arose, they did not always know how to respond.

Fortunately, many of the automated aspects of the process made it possible for less-experienced personnel to be trained to perform some key roles. Significant training time was still required, including many “shadow” shifts giving on-the-job exposure to the process. Leadership positions on the team were filled from within the team, reducing the training time required. Formal on-line training materials were produced to aid the incorporation of new personnel. (Such materials were not required during prime mission, since all personnel were familiar with the spacecraft and their roles.)

VIII. Other evolutions

An additional consequence of the continued operation of the MER rovers is that the content of each planning cycle, rather than becoming simpler, is actually becoming more complex, due to the factors described below. As a result, the length of the tactical shift duration is continually tracked, and changes are incorporated—to the process, the tools, training materials, instructions to the team leads, new interfaces with other spacecraft, or other areas—to keep improving the tactical operations process, to keep it manageable for the team, and to maximize the science return of our rovers while they continue to amaze us with their longevity.

A. Aging rovers

As the rovers continue to operate well beyond their original expected lifetime, inevitably components begin to fail. Upon landing, one of the rovers had an arm actuator heater stuck on, necessitating a change to how that rover was operated, in order to preserve energy for science activities. In addition, over the last nearly 1600 sols of combined rover operations, various robotic actuators have exhibited flaky behavior or have failed outright. In response, the team has devised workarounds and new operation techniques to continue to maximize mobility and science return. These new techniques, however, significantly increase the complexity of both planning and generating the command load.

B. New Flight Software

MER has had two upgrades to its on-board flight software since landing on Mars, with a third planned for the summer of 2006. The new flight software has often incorporated fixes which help to reduce operations complexity in some areas, but also added capabilities, the use of which—while augmenting the rovers’ ability to continue operating, to increase traverse capability, and to gather and return more science data—often increased operations complexity beyond the level originally experienced.

C. Changing Martian Seasons

The MER rovers have also now operated through more than an entire martian year, encountering every martian season. As of this writing, the MER rovers, both of which are in the southern hemisphere, are experiencing martian autumn and are rapidly approaching winter. The changing seasons each bring challenges, which the operations strategy must adapt to: during spring and summer, solar energy is increasingly plentiful, but care must be taken so the rovers and their instruments do not overheat. During fall and winter, the sun drops closer to the horizon, significantly decreasing the solar energy available for rover operation, and heaters must be used more often to keep the rovers and their instruments warm enough, taxing the available energy further. In addition, atmospheric opacity due to dust changes with the seasons, affecting available energy, and the “cleaning events” which occasionally fortuitously remove dust from the solar panels may be tied to particular seasons. Changes are made to the tools to aid manageable operations with each change in season.

D. Changing Operations Environment at Mars

The longevity of the rovers has meant that they are still operating even as new spacecraft enter orbit or attempt to land at Mars. The challenges of these changes include increased scarcity of communications opportunities with the rovers, and increased coordination of MER operations with the operations teams of other spacecraft, whose tactical processes have longer latencies and are out-of-sync with MER’s 5-day-a-week, modified Earth-time schedule. In response to unavailable uplink opportunities, the team has developed two strategies: where possible, the team uses multi-sol plans to make up for sols without uplink opportunities. If the period without uplink is not supportable in this manner, the team developed the capability to command via relay through the Odyssey orbiter—built into the spacecraft during development, but not exercised except as a brief test until the third extended mission, when operations tools and relay coordination processes with the Odyssey operations team were developed to adapt the existing MER tactical process to the requirements of relay commanding.

IX. Conclusion

Even after more than two years of Mars surface exploration, the MER operations design remains in motion. Changing mission resources, new technology experiments, and the challenges of aging hardware continue to drive improvements and adaptations. This operations evolution will likely go on for as long as Spirit and Opportunity survive.

Acknowledgments

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

Andrew H. Mishkin, Daniel Limonadi, Sharon L. Laubach, and Deborah S. Bass, "Working the Martian Night Shift: The MER Surface Operations Process," to appear in IEEE Robotics Automation Magazine Special Issue on Mars Exploration Rovers, June 2006.