Implementing Distributed Operations: A Comparison of Two Deep Space Missions

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

Two very different deep space exploration missions—Mars Exploration Rover and Cassini—have made use of distributed operations for their science teams. In the case of MER, the distributed operations capability was implemented only after the prime mission was completed, as the rovers continued to operate well in excess of their expected mission lifetimes; Cassini, designed for a mission of more than ten years, had planned for distributed operations from its inception. The rapid command turnaround timeline of MER, as well as many of the operations features implemented to support it, have proven to be conducive to distributed operations. These features include: a single science team leader during the tactical operations timeline, highly integrated science and engineering teams, processes and file structures designed to permit multiple team members to work in parallel to deliver sequencing products, web-based spacecraft status and planning reports for team-wide access, and near-elimination of paper products from the operations process. Additionally, MER has benefited from the initial co-location of its entire operations team, and from having a single Principal Investigator, while Cassini operations have had to reconcile multiple science teams distributed from before launch. Cassini has faced greater challenges in implementing effective distributed operations. Because extensive early planning is required to capture science opportunities on its tour and because sequence development takes significantly longer than sequence execution, multiple teams are contributing to multiple sequences concurrently. The complexity of integrating inputs from multiple teams is exacerbated by spacecraft operability issues and resource contention among the teams, each of which has their own Principal Investigator. Finally, much of the technology that MER has exploited to facilitate distributed operations was not available when the Cassini ground system was designed, although later adoption of web-based and telecommunication tools has been critical to the success of Cassini operations.

Nomenclature

JPL	=	Jet Propulsion Laboratory
MER	=	Mars Exploration Rover
OSS	=	Operations Storage Server
Pancam	=	Panoramic Camera
sol	=	1 Martian day
SOWG	=	Science Operations Working Group

I. Introduction

BOTH the Mars Exploration Rover (MER) and Cassini missions have implemented distributed science operations. The MER rovers, intended to perform both in situ and remote observations at two distinct sites on the surface of Mars, were each originally assumed to have a 90-sol (Martian day) surface lifetime. To handle inherent uncertainties in the rovers' interactions with the Martian terrain, the MER surface operations process was designed to be highly reactive, producing new rover activity plans for a given Martian day in response to the results realized from the prior plan's execution. Given the expectation of a short intensive surface mission, all personnel were co-located at JPL and no distributed capability was implemented for the prime mission. However, as it became clear

[*] Principal Engineer, Planning and Execution Systems Section, MS 301-250D
[ε] Cassini Science Planner, Planning and Execution Systems Section, MS 230-205

that the rovers would continue to operate for an extended period, the MER process evolved to enable sustained lower cost operations, eventually resulting in a significant reduction in operations team size and the return of the vast majority of science team participants to their home institutions, from which they continued to support the tactical operations process remotely. The MER mission has now been operating with a distributed team approach for a longer period (> 1 year) than with its original co-located team design.

For Cassini, in contrast, a distributed mission operations system was envisioned from the time of the announcement of opportunity. With a seven-year cruise phase, four-year prime mission and early expectation of extended mission, absence of scientists from their home institutions for the duration of the mission was not a viable approach. Also, locating operational responsibility for the instruments near experts on the instrument and its goals offered possibilities for improved application of scientific and technical judgment to each observation and for reduced costs. Although mission designers recognized that lack of margin in the design of spacecraft resources and that the extent of subsystem interaction would complicate operations, their expectation was that instrument operators would need limited knowledge of spacecraft operations and of details of the ground system. In reality, the complexity of resource allocation and of pointing design, especially after elimination of the scan platform, has meant the deep involvement of the remote participants in virtually all aspects of Cassini operations. A key factor in the success realized by the Cassini mission to date has been the willingness of the participants to step up to the challenges of this operational complexity.

While both MER and Cassini are robotic planetary exploration missions, they operate under vastly different planning timelines (hours vs. weeks), with very different resource constraints, and in radically different environments (planetary surface vs. Saturnian system). We compare the drivers behind the choice to distribute the science teams for the two missions, and consider the challenges associated with enabling effective operations with such distributed teams. We examine similarities and differences between the distributed operations process implementations and experiences for the two missions, and identify lessons learned that may be applicable to the design of operations for future planetary exploration missions.

II. MER Mission

A. MER Mission Overview

The MER mission was designed to land two robotic rovers at distinct sites on the surface of Mars, with the intention of investigating the geologic history, and searching for evidence of the past presence of liquid water and therefore the potential for past life to have existed on the planet. The first MER rover, named “Spirit,” was launched on June 10, 2003, followed by the launch of “Opportunity” on July 7. Spirit and Opportunity landed on opposite sides of the Red Planet on January 3 and 24, 2004, respectively. Spirit began the exploration of Gusev Crater, postulated to have once be a crater lake, while Opportunity roamed through Meridiani Planum. Although neither vehicle was originally anticipated to traverse farther than about 600 meters, or survive significantly longer than its 3-month prime mission, over the next two Earth years, the combined traverse distance of the two rovers has exceeded 10 kilometers.

The rovers are identical in design and instrument complement. The 5 degree-of-freedom Instrument Deployment Device (IDD) has approximately the same reach as a human arm, and can place its turret-mounted instruments—a Moessbauer spectrometer, Alpha Particle X-ray spectrometer, and microscopic imager against or near in-situ targets. In addition, the turret contains a Rock Abrasion Tool (RAT) capable of grinding away the first few millimeters of rock surface, removing the weathering rind to reveal potentially pristine material beneath. The rover instrument complement also includes mast-mounted instruments: the multi-spectral high resolution stereo Panoramic Camera (Pancam), and the Mini-Thermal Emission Spectrometer (Mini-TES). Three additional stereo camera pairs support vehicle traverse planning, autonomous navigation, and operation of the IDD.

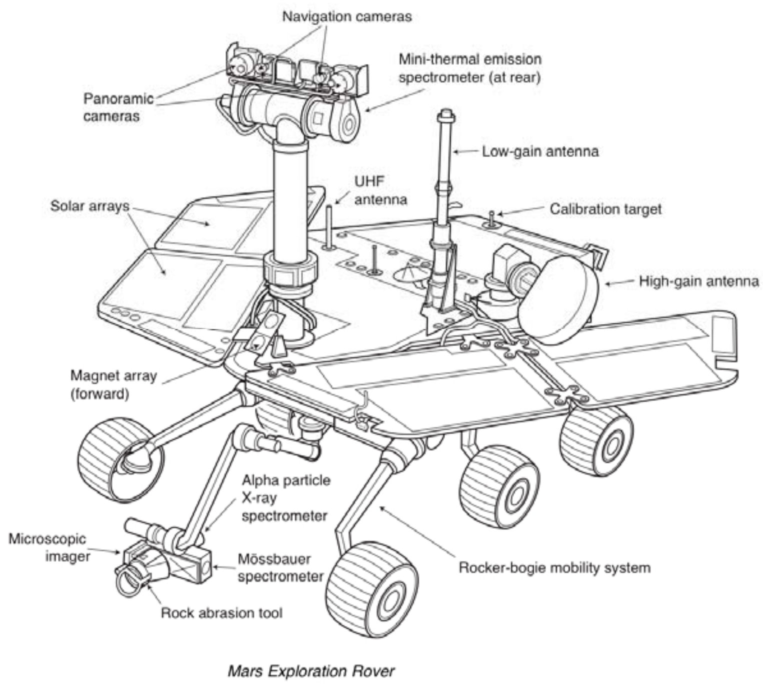

Figure 1. MER Rover Features

B. Initial Operations Design

The design of MER operations responded to two key assumptions: 1) the rovers were wasting assets with limited lifetimes, and, 2) new rover activity plans could not be generated until the results of prior activity plan execution were known by the operations team. The rover lifetime constraint was based in large part on the expectation that dust accumulation on the solar arrays would in time leave the rovers with too little energy to make observations, or eventually, to even support overnight survival heating. The telemetry dependency of rover plans was a consequence of the nature of surface missions: science targets would be selected based on what could be observed from the rover from its current location, and the results of rover traverse and instrument placement activities could not be precisely predicted, due to slippages and other terrain interactions.

MER operations was dominated by the tactical "overnight" timeline which supported daily commanding of the two rovers (see Fig. 2)[1]. The tactical process proceeded from receipt of telemetry from the rover at the end of its day, through assessment of rover state, design of a preliminary science activity plan, refinement and validation of an integrated science and engineering plan, command sequencing, integration, and uplink to the rover in roughly18 hours. The product of the planning cycle was a set of stored command sequences that would govern the rover's activities for a full Martian day. In order to stay synchronized with the solar-powered rovers, the operations team worked on a sliding Mars-time schedule, starting its shift 40 minutes later each day. This Mars-time schedule made available to the team the maximum number of workhours between the downlink in the Martian afternoon and the uplink on the next Martian morning. Since Spirit and Opportunity were on opposite sides of Mars, the tactical process was implemented twice, with two independent tactical teams effectively living in different time zones, twelve Mars hours apart. All personnel participating in the tactical process were co-located in the MER Mission Support area at JPL.

Prior to the afternoon downlink, the science team would meet to discuss the rough science plan for the rover's upcoming sol, based on the expected successful execution of the current sol's plan. Immediately after receipt of downlinked telemetry, the engineering and science teams would assess the state and health of the rover and its instruments, respectively. Over the next few hours, members of science theme groups would assess imagery and other data received, and begin constructing observation plans for the next sol. At the two hour Science Operations Working Group (SOWG) meeting, competing activity plans would be merged and pruned. The leader of the meeting (who was also the science team leader for the entire workshift) was the SOWG Chair. Attendees at the meeting included science theme group members, instrument sequencers, and the uplink engineering team. The team employs software tools during the meeting to aid activity planning, resource modeling, and visualization. The activity plan

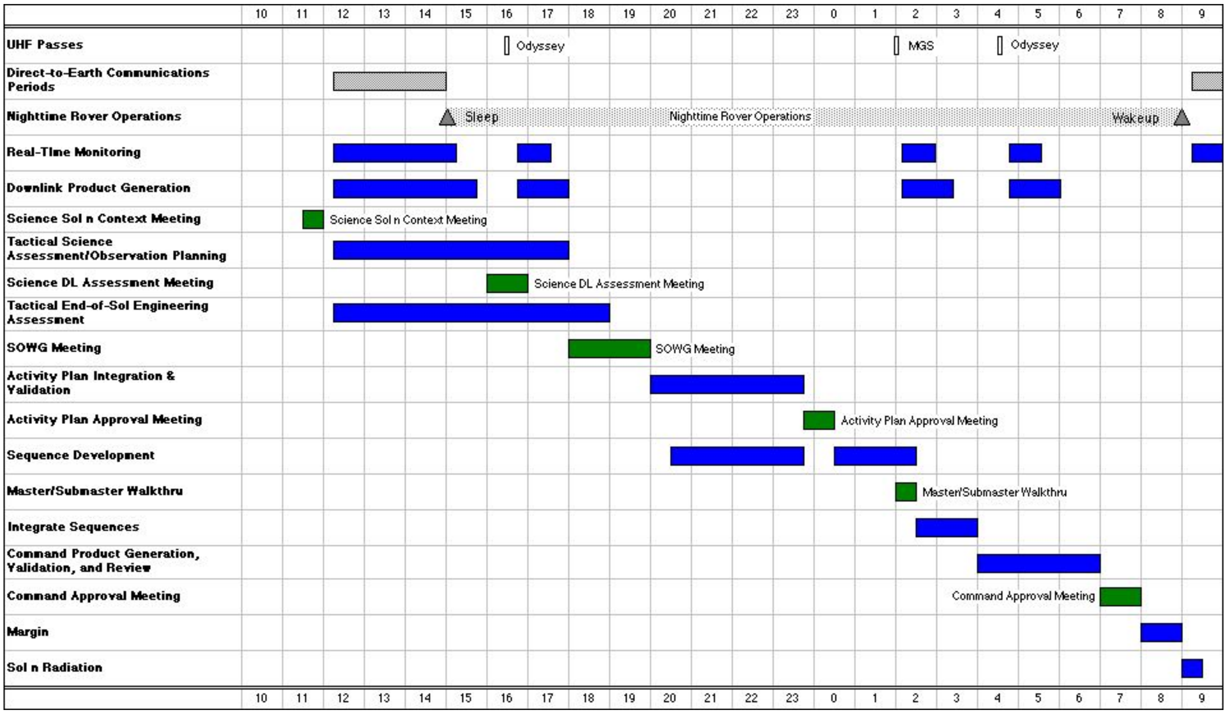

Figure 2. MER Tactical Timeline.
Times shown are in Mars hours, and represent Mars local time at the rover landing site.

under development was modeled for resource usage (time, energy, data volume), and evaluated both from the standpoint of the rover's capacity to execute the plan, and the sequencing team's capacity to implement the plan in the remaining hours before uplink. The team prioritizes the activities in the plan from two perspectives: importance of including the activity in the plan, and the importance of the downlinking the data generated by execution of the activity to the ground. The plan is intentionally over-subscribed by up to twenty percent. The SOWG meeting ends with the electronic delivery of the preliminary science activity plan for the sol.

At the end of the SOWG meeting, most of the science team members are freed to return to science analysis. The subset of the team responsible for delivering the command load then assembles in a single sequencing room, co-located to facilitate efficient execution of the next steps of the process. This team refines the activity plan, incorporating engineering requests, scheduling activities precisely, and performing high-fidelity energy, time, and data volume modeling. Activities that cannot be made to fit are deleted from the plan. (There is no time to re-assemble the full science team and debate such modifications to the plan. Instead, the activity priorities specified during the SOWG meeting guide the choice of items to be removed. The SOWG Chair is present to resolve intra-science trades.) While the activity plan is being refined and validated, other team members are building the individual command sequences that will implement the individual activities. Some of these sequences, for activities that are eventually deleted from the plan, may never be used; but this risk is small, compared to the cost in time of waiting for activity planning to be completed.

At the Activity Plan Approval Meeting, the team reviews and approves the final activity plan and rover motion plan. Team members turn from their workstations located around the periphery of the sequencing room, and face the projection screen at the front of the room to participate in the meeting. The agenda for the meeting is an automatically-generated html document, hyper-linked to all files to be reviewed. This mechanism facilitates the review process, and helps ensure that all items are addressed. Individual workstations can be selected for display on the projection screen, enabling team members to present active files and rover motion simulations.

Sequencing continues for another two hours, at which time the master and submaster sequences implementing the activity plan is reviewed. These sequences govern when the rover wakes ups, shuts down, and triggers the sequences executing each individual activity. All team members must have delivered and validated their sequences by this point.

Over the next 4.5 hours, the full set of sequences comprising the command load is integrated, and all command files and review products are generated. The command load is validated by a combination of manual and automated flight rule checks and software simulation.

Finally, the team reviews the command load at the Command Approval Meeting, verifying that all necessary checks have been completed and open items addressed. Baring uncorrected errors identified during the validation steps, the command load is approved and delivered for upload to the rover by 0900 Mars local time.

C. Evolution to Distributed Operations Paradigm

Spirit and Opportunity were intended to complete their Mars surface missions in 90 sols, as measured from each rover's landing day. (In fact, all mission success criteria—distance traveled, number of distinct locations investigated, and number and types of observations completed—were satisfied within this time period.) But as it became clear that both rovers would be capable of continued exploration beyond the 90-sol plan, the project began working toward modifying the operations design for sustainability. The mission, originally perceived as a sprint, would now become a marathon. No longer was the overriding consideration extracting the last possible bit of useful science return from the rovers before they expired; instead, operations would need to be sustained indefinitely without burning out the operations team, which would in the long run provide greater science return by reducing the probability that human error would put the mission at risk.

The first step was to move from a Mars time staffing schedule to a more human-friendly "Earth-time" schedule. During the first two months of surface operations, the duration of the tactical timeline had dropped to about 11 hours. This reduction was due to a combination of factors, including accumulation of team experience, continuing automation of previously manual steps, and buildup of command sequence libraries. To enable a more regular schedule, the time margin made available by the reduced timeline duration was utilized to eliminate the midnight shift that would otherwise be required. The team used multiple strategies to handle the challenges of the rover downlinks marching through the Earth workday. One of these strategies is to rely on day-old rover status information when more current data will not be downlinked until too late in the Earth day; use of stale data results in some restrictions on safe rover activities, for example permitting rover driving only every other day until the Earth-Mars time phasing becomes more favorable.

Many members MER science team at large, who were involved in science analysis and observation planning, but not the later steps of the tactical timeline, began attending their key meetings via telecon, including the SOWG meeting, soon after the end of prime mission. By September 2004, even SOWG Chairs began to participate via telecon and videocon, with the displays of key software tools showing activity plans and rover motions provided to them by webcam. Around this time, developers of command sequences for science instruments began to produce these sequences at their home institutions using the same software tools they had used at JPL, delivering their sequences to the same locations they had previously. The design philosophy for distributed operations was one of virtual co-location: team members who were not physically present would be continually available and in contact during the tactical shift via an open teleconference line. While virtually all of the science team participates in MER operations remotely, engineering team members, mission managers, and project management all continue to be co-located in the JPL MER operations facility. Since the engineering team is comprised of JPL employees, there has been no drive to distribute this team or its functions.

III. Cassini Mission

A. Cassini Mission Overview

Figure 3. The Cassini Spacecraft with the Huygens Probe descending to Titan and Saturn in the background

Cassini-Huygens is dual mission consisting of an orbiter conducting a tour of the vast Saturnian system and a probe, which descended through the Titan atmosphere for the first landing on a surface in the outer Solar System. After successful launch on October 15, 1997, the spacecraft spent nearly seven years getting to Saturn using four gravity assist maneuvers (Venus-Venus-Earth-Jupiter). During the long cruise phase, limited science was performed to checkout and characterize the instruments. Further preparation came from science collection during the months of the Jupiter flyby, with closest approach on December 30, 2000. Shortly before arrival, the intensity of science observations stepped up with observations of Phoebe on June 11, 2004. During the critical Saturn Orbit Insertion (SOI) on 1 July 2004, Cassini made its closest approach to the planet's surface of the entire mission at an altitude of only 0.3 Saturn radii (18,000 km) and crossed through Saturn's ring plane at a gap in the thin outermost area. It then commenced a tour with 76 orbits around Saturn which include 44 close Titan flybys and 8 targeted icy satellite flybys. On December 25, 2004 the probe separated from the orbiter finally reaching Titan twenty days later. Descending via parachute, it sampled the atmosphere and provided the first and closest view of its surface. Primary science objectives for the mission include (1) close-up studies of Saturn composition and atmosphere as well as its interesting interior and magnetic environment (2) Saturn's spectacular rings (3) its moons especially the variety of unique surface features (and emissions from Enceladus) that characterize its icy satellites, (4) the intriguing moon Titan, Saturn's largest, with its brownish-orange, dense hazy methane-rich atmosphere of nitrogen that may provide clues to the chemistry of primordial Earth. The prime mission ends in mid-2008, for a total mission duration of 10.7 years. Planning is now beginning for a two year extended mission.

Cassini-Huygens is the largest interplanetary spacecraft ever built and one of the most complex. At launch, the spacecraft had a mass of about 5600 kilograms. Cassini is a three-axis stabilized spacecraft whose components include one high gain and two low gain antennas, two solid state recorders, three Radioisotope Thermal Generators (RTGs) for power, and an attitude control system with a main engine, thrusters, and reaction wheels.

Cassini's twelve science instruments are grouped into three categories: Optical Remote Sensing, Fields/Particles/Waves and Microwave Remote Sensing. The instruments are all mounted to the body of the spacecraft. The optical remote sensing instruments are roughly co-aligned so they can sometimes collect data collaboratively. With suitable choice of the secondary pointing axis, the MAPS (Magnetosphere & Plasma Science) instruments can also be simultaneously observing.

Table 1. Cassini Instrument Teams

INSTRUMENT	SCIENCE	PRINCIPAL LOCATION
Cassini Plasma Spectrometer (CAPS)	measures the energy and electrical charge of particles such as electrons and protons	Southwest Research Institute
Cosmic Dust Analyzer (CDA)	measures the size, speed, and direction of dust grains	Max Planck Institute
Composite Infrared Spectrometer (CIRS)	measures infrared energy to study temperature and composition	Goddard
Ion and Neutral Mass Spectrometer (INMS)	examines neutral and charged particles in extended atmospheres and ionospheres	U Michigan
Imaging Science Subsystem (ISS)	captures images in visible, near-ultraviolet and near-infrared light	Space Science Institute
Dual Technique Magnetometer (MAG)	measures the strength and direction of the magnetic field	Imperial College
Magnetospheric Imaging Instrument (MIMI)	images Saturn's magnetosphere and measures interaction with the solar wind	APL
Radio Detection and Ranging Instrument (RADAR)	radar imager maps of Titan's surface and measures the height of surface objects	JPL
Radio and Plasma Wave Science instrument (RPWS)	measure the electric and magnetic wave fields	U Iowa
Radio Science Subsystem (RSS)	measures telling changes in radio waves sent from the spacecraft to antennas on Earth	JPL
Ultraviolet Imaging Spectrograph (UVIS)	measures ultraviolet energy to study structure, chemistry and composition	Laboratory for Atmospheric and Space Physics, U Colorado
Visible and Infrared Mapping Spectrometer (VIMS)	captures images using visible and infrared light to identify chemical compositions	U Arizona

B. Operations Design

Motivated by the desire to keep science operations near the expertise on the instrument and its capabilities, the Cassini-Huygens designers envisioned remote science operations from the beginning. The long mission duration—nearly eleven years from launch to end of prime mission—meant that it was impractical to relocate the scientists from their home institution to JPL. The next section describes how current operations are different from this initial vision. This section details how distributed uplink operations are actually accomplished for Cassini.

Achievement of science objectives began with selection of the tour, highlights of which were described above. Various types of opportunities (e.g. flybys, occultations, Saturn periapse) were identified within the selected tour. Having the entire science community address dividing the entire tour among the various instruments and objectives proved unworkable, so the tour was divided into roughly 200 segments each of which was assigned to one of six science discipline working group (e.g. rings, Saturn, Titan). These groups met weekly via teleconference for over three years to integrate the tour into a conflict-free timeline of activities. The decision makers were scientists from around the world supported by science operations personnel both at JPL and at remote sites, by science planners at JPL and by spacecraft team members. Some scientists, particularly those from small teams, participated in multiple groups. They reached preliminary agreement on all negotiated resources, specifically pointing time, collaborative data collection, data volume, telemetry mode, and power mode.

Most of this activity plan is accomplished by executing autonomous sequences stored onboard the spacecraft. These sequences include both spacecraft commands—at the system and subsystem level—and instrument commands. Construction of the sequences is accomplished from a dedicated mission support area at JPL and from science instrument operation centers and sub-centers geographically distributed throughout the United States and Western Europe. Functions centralized at JPL for the scientists include mission planning; sequence integration and system level validation; uplink radiation through the DSN; telemetry data collection, basic processing, and storage; spacecraft monitoring; and spacecraft navigation.

After integration, an initial implementation of this plan to the command level was completed for each sequence in the tour. The effort, which took 2.5 years, was coordinated by science planners at JPL but involved extensive participation from the instrument operation teams in supplying commands and validating and correcting the commanding. At any one time, multiple sequences were in various phases of this effort. The logistics of distributed sequence generation were daunting at first. Teleconference meetings were initially required multiple times per week, document availability and exchange required web posting and ftp sites, email was extensively used to communicate and document communication, and phone bills were considerable.

At 20 weeks before execution of the sequence begins, a window for aftermarket changes to the implemented product is opened. These changes may include responses to new discoveries or adjustments for instrument performance changes. Approval of changes is negotiated by the discipline working group that originally integrated the sequence. The plan is then re-verified as conflict-free and re-implemented to the command level.

Under the leadership of the uplink operations office, the sequence virtual team verifies and validates all commands to be sent to the spacecraft, although instrument internal commands are not included in the background sequence. Any problems in commands from the distributed sites are returned to the site for resolution. The sequence virtual team is also responsible for all system-level real-time commands associated with the sequence, and monitors the progress of the activities until the sequence is complete. Instrument teams are vitally involved in the uplink process wherever they are located in order to ensure that their observations are correctly implemented.

Changes once the sequence is executing are limited. An instrument's internal commands may be sent via the instrument-internal real-time command process. Updates to the vectors that control pointing may also be uplinked if improved knowledge of the spacecraft or target ephemeris is available.

C. Evolution of Distributed Operations Paradigm

In the original definition of distributed operations, a limited set of resource allocation templates called operational modes was to be defined. The Project Science Group would use these templates to construct an activity timeline. The instruments would then be free to command within the constraints of the applicable mode with only limited interaction with engineering personnel at JPL. This model bears only limited resemblance to the actual operations concepts.

With the exception of power, predefined patterns of resource allocation proved too confining. Also, concentrating the decision making on opportunity selection and resource allocation in the Project Science Group proved unwieldy and too time-consuming. These problems led to the dispersed science groups dealing more flexibly with opportunity assignments and resource allotment as described above.

Perhaps the most drastic influence on the operational paradigm was the elimination of the instrument scan platform in a move to cut spacecraft costs. The instruments were then body-mounted so that pointing an instrument requires pointing the entire spacecraft, a change with profound ramifications (and additional cost) for operations. Since the daily nine-hour downlink requires pointing the antenna at earth, the observation pointing time was immediately reduced by more than 1/3 increasing pressure on the planners to optimize available time. Science pointing design was also significantly complicated since in addition to pointing the instrument as desired it is necessary to avoid turning through or pointing at any geometric hazard (e.g. bright body in stellar reference unit FOV, heating on infrared coolers). To mitigate the increased pointing design workload, the decision was made to distribute pointing design to the observing instrument. This eroded the barrier isolating the science operations teams from engineering so that the Cassini paradigm was not longer distributed science operations but distributed operations overall.

IV. Comparison of MER and Cassini

A. Key Similar and Distinct Driving Constraints and Requirements

1. Time Scale

The MER operations paradigm was to generate a new command for each Martian sol, or once every 24.6 hours. Each command load governs the rover's activities for one Martian sol, with "runout" sequences to perform generic science observation for an additional 1.5 sols in the event that the next command load is not received onboard. Early in prime mission, the tactical command-turnaround process required about 18 hours, leaving little margin in the planning cycle. Over time, various improvements have reduced the planning cycle to about 8 hours.

Some MER operations processes take place on a more strategic schedule; by definition, any process requiring more time than the tactical process allows is "strategic." Strategic activities include negotiation of communication opportunities with other missions and the DSN, testbed validation of first-time or risky rover activities, and planning of long-term science campaigns.

For commanding purposes, the Cassini four-year tour is divided into 41 sequences with a mean duration of 37 days. Working from a provisionally implemented activity plan, the sequence development process takes twenty weeks. Practically, this means that five sequences are somewhere in the process from kick-off to execution concurrently. Most personnel, both scientists and engineers, are supporting activities for multiple sequences simultaneously.

Strategic redirection, as in response to new discoveries, thus is typically at least six months downstream of expressed desire. In most cases, because of both the level of effort required to redo the plan and of the length of time required to generate commanding, new activity is targeted for the extended mission.

Tactical alteration, particularly at the instrument level, is available on a limited scale through real-time commanding. Some adjustments, such as pointing updates, are anticipated and built into the plan. Immediate response to telemetry data is largely impossible because of the light-time delay (1 hour 10 minutes ± ~10 minutes), because there is no communication with the spacecraft except during planned downlinks, and because of potential consequences of interfering with the sequence commands.

2. Complexity/Scope of Command Load

A single MER sequence plan may consist of 500 to 1000 commands, governing the rover's activities for one Martian sol, plus 1.5 to 3 sols of "runout" sequence. The runout sequence specifies what actions the rover should take in the event that the next sol's command load is not received as scheduled, and generally includes generic remote science observations to be performed in the absence of more informed instructions received from the ground. The command load for a given sol will usually consist of fewer commands than the corresponding sequence plan, since the sequence plan often incorporates reusable subsequences that are already onboard, and which will therefore not be included in the command load to be uplinked to the spacecraft.

A Cassini background sequence, containing commanding for a period of 25 to 44 days, must fit into the approximately 100K words available in spacecraft memory. The sequence is constructed from 1250 commands stems and must satisfy 305 flight rules as well as other constraints. During each sequence, from 21000 to 46000 commands are executed from the background sequence, maneuvers, mini-sequences, and real-time commands (but not including instrument internal commanding). The daily background command load varies from as low as 125 to as many as 10,000 depending on the intensity of activity and the number of commands required for particular activities.

3. Science Team Organization

The MER mission has a Project Scientist and Deputy Project Scientist to lead the science team. However, since the entire rover instrument suite is under the cognizance of a single Principal Investigator, it is the PI who is most involved in the management of the science team. Payload Element Leads have responsibility for individual instruments, including managing and scheduling the specialized science team members who assess the health of their respective instruments and generate the instrument command sequences. Tactically, there is an SOWG Chair for each rover, who leads the science team during the execution of the tactical process, adjudicating intra-science disputes and priorities. However, since a clear chain of command is absolutely essential during the tactical timeline to ensure on-time delivery of the command load, the conduct of the tactical process is ultimately under the direction of the Tactical Uplink Lead, an engineering team member who ensures that science and engineering needs are appropriately integrated, that spacecraft health and safety concerns are addressed, and that the schedule is maintained.

Within the science team are four Science Theme Groups (Atmospheric Sciences, Geology, Geochemistry/ Mineralogy, and Soils and Physical Properties) that each review telemetry and propose observations relevant to their themes. A first theme group—Long Term Planning—acts a bridge between tactical and strategic processes, looking several sols ahead to ensure continuity of the science plan. The only science team role that is in almost all cases performed on-site at JPL is the so-called Keeper of the Plan, a role that became necessary as distributed operations became pervasive on MER; the Keeper of the Plan integrates activity requests from the remote theme groups, operates the science activity planning tool during the SOWG meeting, and ensures delivery of the preliminary plan to the engineering team at the conclusion of the SOWG meeting.

The Cassini science organization is headed by a Project Scientist and Deputy Project Scientist at JPL. Each of the twelve instruments and the Huygens Probe have a science team composed of a Principal Investigator (Team Lead for facility instruments), Co-Investigators, an Investigation Scientist at JPL, research assistants and graduate students, an Operations Team Lead, and operations engineers. There is also a cross-discipline group of scientists not associated with a particular instrument. From science team to science team, the role definition is not uniform. On some teams, scientists are active in observation design, uplink implementation, instrument monitoring and data retrieval. On other teams, these functions are assigned to dedicated operations engineers.

Each science team has a main site with JPL furnished hardware and software. Some of the larger teams have multiple sub-sites with scientists and operations personnel who participate in science operations through both direct interaction in the implementation process and interaction with operations personnel at the team's main site. These sites may run project supplied software on hardware supplied by their institution or their science team.

4. Predictability of Spacecraft State

For MER operations, responding to new information on a daily basis in the norm. While strategic science targets (e.g., craters) can be identified weeks or even months in advance based on orbital imagery, plans for the next sol depend on data collected by the rover on the current sol. Energy available for rover activities will be determined by the attitude of the solar-powered rover at the end of its drive; the next science target will likely be determined by what is in view or in reach in images from the rover's current vantage point. Wheel slippage during traverse and the need to autonomously avoid hazards contribute to uncertainty in vehicle position until end-of-drive images have been downlinked to the operations team.

For Cassini, in contrast, the selection of science opportunity is largely driven by known opportunities in the tour design. Most of the factors in the environment that impact spacecraft state (e.g. gravity) are well-enough understood that only minor adjustments are required to compensate. Surprises, even when representing science discovery, are somewhat unwelcome because of the workload associated with replanning. An example of such an adjustment would be the need to raise the altitude of the Titan flybys because the atmosphere is denser than anticipated and might result in undesirable torque to the spacecraft. Adjusting the tour to compensate required many work-months of effort and considerable iteration including significant participation from the distributed scientists.

5. Contention for Resources

Resource allocation among competing science requests must be resolved within the first few hours of the MER tactical process. The basic type of sol (e.g., traverse, in situ science, or remote science) is tentatively determined in early science meetings, or may be clear from the activities of the prior sol; this context bounds the activities for the upcoming sol. By the start of the SOWG meeting, the engineering team has provided the resource constraints for the next sol, including energy, downlink data volume, available onboard data storage, and duration of the rover's day. MER-developed software tools provide early rough estimates of the resources required for all proposed science activities. The SOWG Chair must guide the science team to a preliminary science activity plan, with activity priorities, by the completion of the SOWG meeting.

One of the key complicating factors in distributing Cassini operations is that the scientists and their operations teams are intimately involved in negotiating contested spacecraft resources. While Cassini is a highly capable spacecraft, key resources are insufficient to support all desired use simultaneously. Cassini does not restrict science collection to one instrument at a time. At periods of high opportunity (e.g. periapse or fly-by) it is common for all of the ORS instruments and all of the MAPS instruments to be operating simultaneously. The various instruments must make concessions in required power, bus data rates, data volume collected and data storage, thermal constraint satisfaction, and pointing direction. Negotiating these concessions is complicated by the fact that is difficult to get all parties to the compromise connected at the same time. Because so many resources are tightly coupled to the plan, revisions of any sort tend to ripple into reallocation of multiple resources. For example, substitution of a requested DSN station may require altering telemetry modes which will reduce downlink capacity which will require less data storage which will require renegotiation of the timeline.

B. Distribution of Hardware and Software

Workstations, configured with MER-developed activity and command sequence planning and visualization tools, were provided to each of the remote sites as MER moved into its distributed operations configuration. Key tactical planning files are automatically pushed out to the secure remote sites, where local subsets of the MER Operations Storage Sever (described below) are maintained. At the remote sites, partial sol activity plans and instrument sequences are produced, and delivered to the Operations Storage Server at JPL.

Other MER operations functions, and their associated hardware and software, remain located in the JPL Mission Support Area. Among these functions are high fidelity resource modeling, rover motion planning and simulation, sequence integration, software simulation, generation of command products, overall command load validation, and transmission of the command files to the DSN.

The original distributed operations concept for Cassini called for a SOPC (Science Operations and Planning Computer) supplied by JPL to each team's principal operations center to act mainly as a transit port for instruments commands out and telemetry in. This proved to be a radical under-estimation of the connectivity required, but Cassini was fortunate in that rapid improvements in network technology proved adequate. Nonetheless, maintenance of the electronic connections between JPL and the distributed sites requires significant personnel at both.

In the original vision where distributed sites provided instrument commanding, little project uplink software would have been required there. When the distributed sites became involved in providing commands to the background sequence, the whole mission sequence system was needed at the team sites. In addition, downloadable versions of the pointing design tools were required to support scientists not at the principal sites. It was also necessary to replicate the input file structure and update configuration files multiple times per sequence. Training and support had to be provided to remote users. Providing software to run on non-project hardware proved a significant burden to the small development staff as did demands for versions for alternate platforms. These were met to the extent possible with a limited budget.

C. Communications Infrastructure

1. Videocons and Telecons

The use of video-conferencing on MER has evolved with distributed operations. During prime mission, all participating personnel would be physically present at the daily Science Operations Working Group meetings, during which the preliminary science activity plan for the sol would be defined. As the mission progressed, some science team members would participate by telecon. When the SOWG Chair (the meeting leader) began participating remotely, video-conferencing was used to allow the participants at the two primary sites to see each other. An internet camera in the SOWG room was focused on one of the room's projection screens, providing remote participants with views of either the science observation targets or the rover activity plan under development. Very soon, the team dispensed with two-way video-conferencing, as the combination of teleconferencing and views of the planning tool displays were sufficient to the efficient conduct of the meeting. The number of persons physically present during the SOWG meeting has dropped forty or more down to about ten.

For the later steps of MER planning and sequencing, teleconferencing is continuous. For a given rover, the JPL team members co-locate in one room, and remote participants dial into the conference line for the next several hours, and report when they will be unavailable, but otherwise are continuously available for discussions or questions that may occur between scheduled meetings.

Cassini makes extensive use of telecons, although as the processes become more routine email communication and web postings have replaced some meetings. Scheduling of teleconferences is highly constrained because the beginning of the work day at JPL is the end of the work day for the European teams. Video-conferencing was tried during cruise but dropped because of cost and problems associated with immature technology. One of the advantages MER has over Cassini is that most participants had met one another before they started working over the phone. Even after years on the project, many Cassini personnel have never met their colleagues at other sites.

2. Operations Storage Server / Project Database

MER implemented a Operations Storage Server (OSS) for storage of derived telemetry products and uplink planning files necessary to the operations process. Data are stored on the OSS in a formal directory structure, organized by sol number, to ensure that all planning files and intermediate products have a known, unambiguous location. Pre-generated subdirectories provide specific working areas for preliminary science activity plans for individual science subteams, integrated activity plans produced at the SOWG meeting, engineering activity plans, refined and validated rover activity plans, instrument sequencing, integration of all sequencing products, online reports, and feed-forward to the next sol's planning process. As the planning for a given sol progresses, these directories are populated with the intermediate and then final products. Since each successive sol of planning occurs in a new set of directories, the directories become an archive of the planning process each sol, available for reference whenever needed. Remote participants deliver their instrument sequences to the directories allocated for their instruments, together with electronic sequence delivery forms, which specify how the sequence should be managed onboard.

As with MER, the smooth flow of partial products is essential to an orderly process. However, Cassini does not have the same urgency to product availability. At the port in the process where an integrated sequence product is produced, the teams deliver their sub-sequences to the project database. Some teams use the automatic notification capability of the database while others generate their own email notification. In most cases there is at least one day of margin in the schedule in the event of non-delivery by one of the teams but this has rarely been needed.

One of the minor frustrations associated with having so many different individuals contributing commands was ensuring that the uplink tools were correctly configured at all the different locations. This was resolved by replicating the input file system either by mirroring the central version or by download and providing configuration files at each port of each process for each tool.

3. Web Site and Email

A key mechanism for making detailed planning and analysis data available to the team has been the Uplink and Downlink Reports available via a controlled-access web site. These reports are generated for every Martian sol, and include sections for entries by each tactical operations role. Although each report is viewable as an integrated online document, each section may be edited independently of and simultaneously with all other sections. The Downlink Report provides the latest vehicle health status, in summary and by subsystem, and assesses the success achieved in executing the sol's plan. The Uplink Report summarizes the plan for the sol, provides details of the design of all observations and rover motions, and communicates any issues that must be addressed by the next sol's uplink planning team.

Strategic information is generally communicated to the team by email. This includes personnel shift scheduling, communications window planning, and schedule updates.

Given that distributed operations for Cassini turned out to be far more complex that envisioned by its early proponents, it's fair to say that Cassini operations were saved by the explosion of telecommunications and the internet in the late 1990s. Email distribution aliases have been an essential tool for widely disseminating information relating to project information, working groups, tour design and integration, sequence development, day-to-day operations, and tool development. The email threads are also captured as part of the project documentation process. Email is particularly useful given the wide time gap between JPL on the U. S. west coast and science teams in central Europe.

The web has also been a central factor in the success of Cassini operations. For each sequence, the sequence leads at JPL maintain a web page with schedule and contact information, meeting notices and handouts, reference materials and documentation, and verification and validation results. One of the most successful uplink tools is CIMS, the Cassini Information Management System. This data base of all activity requests provides an up-to-date picture of the science plan via a web interface.

4. Meeting Structure and Process Gates

Key MER tactical meetings (all occurring within an 10-hour period during the extended mission) are indicated in green in Fig. 2. The sol Context and Downlink Assessment meetings provide opportunities for the science team to coordinate initial planning for the sol. The SOWG meeting establishes the initial science activity plan for a sol, with feasibility assessment by the engineering team. At the Activity Plan Approval Meeting, the team confirms that the plan is achievable within available resources and meets the original intent of the science team. The Master/Submaster Walkthrough reviews the commands that implement the activity plan and all rover motions. Finally, the team determines that the command load is safe for transmission to the rover at the Command Approval Meeting.

The rhythm of Cassini Uplink Operations has a longer period than MER but is still quite regular. The Project Science Group holds weekly teleconferences and meets at a physical location three to four times per year. With the tour fully integrated, meetings of scientist discipline working groups for opportunity selection and resource negotiation are held on an as needed basis, though these will become more frequent as work on the extended mission accelerates. As personnel become familiar with the operations process, some teleconference meetings have been replaced by posting of meeting packages to the web site with annotations via email. At the start of command development, a kick-off meeting is always held to clarify expectations and assumptions for that sequence, particularly delivery dates for products from the science teams. Thereafter, the essential meetings occur at points in the process where membership shifts or where sequence validity is addressed. Delivery of products to the process via the project data base is confirmed by email. For the background sequence, deadlines are on the order of a day rather than a particular hour and reflect the desire to retain relatively normal work hours despite having operations across many time zones.

D. Lessons Learned

1. MER

The tight tactical timeline required for every-sol commanding, although extremely demanding on the operations team, appears to have facilitated the transition to distributed operations. The original MER tactical process tightly integrated science and engineering throughout, with team members in close proximity and easily available to facilitate meeting deadlines under extremely time-constrained circumstances. The distributed version of MER operations process is structured nearly identically to the original co-located process, with continuous communications among team members maintained by teleconference lines, webcams, and online documentation.

Online and web-based documentation have proven essential, both for an efficient tactical process and effective distributed teams. Even without webcams, remote team members can access hyper-linked meeting agendas and review the same information as the co-located personnel, and therefore participate effectively by telecon in key review meetings. Online reports and the formal archive of data facilitates rapid retrieval of operations-critical information from prior shifts or sols, regardless of the physical location of team members.

The co-location of the full operations team (including both science and engineering elements) during prime mission provided the critical opportunity to mature the operations process and complete the training of the participants. This approach (which was instituted less through forethought than simply as a consequence of the presumed short mission lifetime) avoided compounding the challenges of a new tactical Mars-time operations process and those of distributed operations. When science team members began to participate in operations from their home institutions, they already had an up-close view of the process they were supporting.

2. Cassini

Establishing a distributed operations environment was far more complicated than anticipated with serious issues in modes of interaction, flow of authority, familiarization with processes, communication across time zones, connectivity, synchronization, configuration management and export of technology.

With multiple teams, the loss of efficiency in decision-making by consensus has been aggravated when the decision makers are not co-located.

Issues in spacecraft operability will be compounded by distributed operations. Lack of margin is harder to resolve with distributed decision makers. The more people who have to deal with workarounds for spacecraft design flaws, the harder operations will be.

The advantages of distributed operations for Cassini do not include significant cost savings. The elimination of the scan platform meant that the science teams required significant knowledge of the spacecraft and familiarity with overall operations. This in turn meant additional staff and longer training times. Engineering responsibilities at the science team sites also required duplication of hardware and software infrastructure.

V. Conclusion

The distinct experiences with distributed operations encountered by the MER and Cassini missions to date suggest some preliminary conclusions:

- Tight time constraints on an operations process may actually facilitate team distribution, by forcing team members into a single-minded focus on operations, which may not be achievable in a more traditional operations environment.
- Many of the same tools implemented to foster communication and the dissemination of information across a largely co-located team are directly applicable to the conduct of distributed operations.
- Distributed operations will tend to be more susceptible to weaknesses in the overall process than co-located operations; however, an initial phase of co-located operations may permit these weaknesses to be identified and corrected prior to the shift to distributed operations, as well as contribute to team-building and training.

Acknowledgments

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

References

[1]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.