

# LESSONS LEARNED FROM COORDINATING RELAY ACTIVITIES AT MARS

Roy Gladden<sup>(1)</sup>, Pauline Hwang<sup>(2)</sup>, Bruce Waggoner<sup>(3)</sup>, Bruce McLaughlin<sup>(4)</sup>, Paul Fieseler<sup>(5)</sup>, Reid Thomas<sup>(6)</sup>, Maria Bigwood<sup>(7)</sup>, Paul Herrera<sup>(8)</sup>

Jet Propulsion Laboratory / California Institute of Technology  
4800 Oak Grove Drive, Pasadena, CA, 91109-8099, The United States of America

<sup>(1)</sup> Roy.E.Gladden@jpl.nasa.gov, <sup>(2)</sup> Pauline.P.Hwang@jpl.nasa.gov, <sup>(3)</sup> Bruce.C.Waggoner@jpl.nasa.gov, <sup>(4)</sup> Bruce.A.McLaughlin@jpl.nasa.gov, <sup>(5)</sup> Paul.D.Fieseler@jpl.nasa.gov, <sup>(6)</sup> Reid.C.Thomas@jpl.nasa.gov, <sup>(7)</sup> Maria.Bigwood@jpl.nasa.gov, <sup>(8)</sup> Paul.N.Herrera@jpl.nasa.gov

## 0.0 Abstract

The Mission Management Office at the Jet Propulsion Laboratory was tasked with coordinating the relay of data between multiple spacecraft at Mars in support of the Mars Exploration Rover missions in early 2004. The confluence of three orbiters (Mars Global Surveyor, Mars Odyssey, and Mars Express), two rovers (Spirit and Opportunity), and one lander (Beagle 2) has provided a challenging operational scenario that required careful coordination between missions to provide the necessary support and to avoid potential interference during simultaneous relay sessions. As these coordination efforts progressed, several important lessons were learned that should be applied to future Mars relay activities.

## 1.0 Introduction

Beginning in late 2003, it was intended to have six spacecraft operating at Mars and communicating in a coordinated fashion in a “relay network,” as shown in Figure 1. Beagle 2, operated by Dr. Colin Pilinger of the Open University for the European Space Agency, arrived at Mars in late December 2003, but was lost on entry. At the same time, the Mars Express Orbiter, operated by the European Space Agency (ESA), successfully arrived at Mars and after several weeks settled into its operational orbit. NASA’s Mars Exploration Rover A (MER-A or Spirit) arrived in early January 2004 and MER-B (Opportunity) arrived several weeks later. Additionally, NASA’s two orbiters, Mars Global Surveyor (MGS) and Mars Odyssey,

operated jointly by the Jet Propulsion Laboratory (JPL) and Lockheed Martin Astronautics (LMA), had arrived in earlier years and remained operational in orbit about Mars. With the loss of Beagle 2, the use of Mars Express in the network was reduced.

Processes were developed to facilitate communication between these very diverse spacecraft and organizations. These processes led to great success in coordinating the relay activities not only between the NASA spacecraft, but also with the international partners. This paper describes the lessons that were learned in the development and exercise of these processes that may be applied to future missions.

## 1.1 The Challenges of Coordinating the Mars Relay Network

With five spacecraft operational at the beginning of 2004, a period of unprecedented activity at Mars ensued. Several difficulties presented themselves in

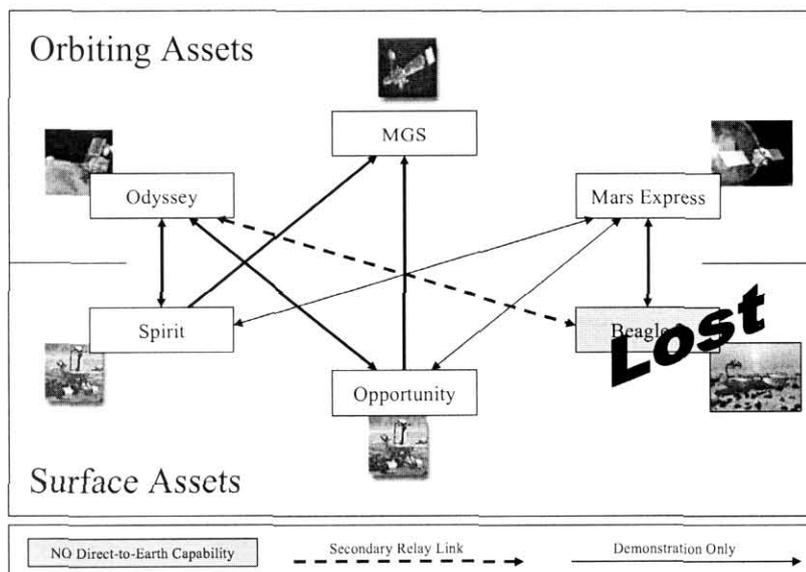


Figure 1: Mars Relay Network in 2004

this relay network that needed to be overcome. Most of the difficulties resulted from three facts:

- 1) The orbiters were<sup>1</sup> not simple data relays; they were extraordinarily complicated vehicles conducting parallel science operations at very low staffing levels.
- 2) Effectively communicating the needs of each project required strong interfaces between the teams, effective training, and thorough testing.
- 3) The various spacecraft in the network were commanded with very different planning horizons, with the lander teams choosing to command their spacecraft on a daily timescale and the orbiter teams operating on a timescale of weeks and up to a month.

For further details on some of the challenges associated with coordinating the relay activities among the various spacecraft, refer to [1].

## 1.2 Background

The process for coordinating relay activities was divided into two sub-processes: a Long-Range Relay Coordination Process and a Short-Range Relay Coordination Process. Additional processes to support the unique commanding needs of each spacecraft were attributable as secondary processes, as illustrated in Figure 2.

The Long-Range Relay Coordination Process was designed to overlay longer-term relay opportunity predictions with estimates of the onboard activities of the relevant orbiters in order to estimate forward-<sup>2</sup> and

<sup>1</sup> While the past tense is used throughout this paper, at the time of writing relay activities were continuing with great success. The orbiters remain healthy and operational and the rovers continue to operate in their extended missions.

<sup>2</sup> “Forward-link data latency” is defined as the duration before the beginning of a particular relay event that a command product intended for a destination asset must be on the project database of the relay asset.

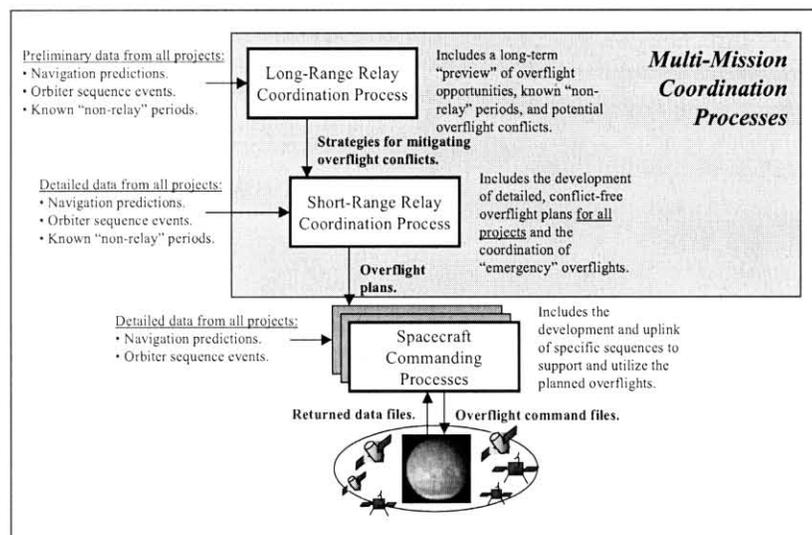


Figure 2: Multi-Mission Relay Coordination Processes

return-link<sup>3</sup> data latencies and to provide a “preview” of the relay opportunities that were available to the rovers. This information was used by the rover teams to determine which overflights were most meaningful for them to utilize for relay purposes.

The Short-Range Relay Coordination Process was the mechanism to determine the details of specific relay opportunities and to resolve any conflicts that existed in the network. This process also produced command sequence files for both Mars Odyssey and MGS and other support files.

There were two primary components of each of the Long- and the Short-Range Relay Coordination Processes:

- 1) Software tools were utilized to calculate the forward- and return-link data latencies, determine the details of each spacecraft’s relay plans, and to check for conflicts in the network.
- 2) Meetings were held where representatives from all the spacecraft in the network came together to resolve issues and to discuss the details of the relay plan.

The Spacecraft Commanding Processes shown in Figure 2 included all processes that were required by each spacecraft to generate the commands necessary to

<sup>3</sup> “Return-link data latency” is defined as the duration after the end of a particular relay event that a “return” link data product should begin being received by an Earth-bound ground station.

enable the relay activities. They were different for each spacecraft and so will not be detailed herein.

## 2.0 Lessons Learned

### 2.1 Mis-Matched Planning Horizons

Both Odyssey and MGS had long-term planning horizons which governed the generation of command sequences to operate the spacecraft. This means that at the beginning of the development cycle, activities were planned to occur on the spacecraft as far as eight weeks in the future. This advanced planning was possible because of the repeating nature of the housekeeping activities and the fact that their stable orbits allowed the planning of science targets well in advance.

For the MER rovers, the planning horizon was much shorter; command sequences were submitted to these rovers as often as several times per day. The objectives on a given day were controlled by the amount of activities completed on the previous day and the changing science objectives created by the daily movement of the rovers across the Martian terrain.

Merging these planning horizons together required some concessions from all parties. First, the orbiters needed to provide a mechanism to perform relay activities on a shorter timescale, and, conversely, the rovers needed to provide a mechanism to plan lander activities on a longer timescale. The initial compromise was for each orbiter to provide weekly sequences in which the lander teams specified parameters for each overflight they were planning to use. Other geometrically feasible overflights were scheduled with default settings for the relay link.

This process worked and continues to function at the time of this writing. Nevertheless, several months before Spirit arrived at Mars, the MER project asked for additional flexibility at a "tactical," or daily level. This was impracticable for the orbiter teams because they could not adapt to producing sequences at this frequency without significantly increasing staffing levels. For Odyssey, one solution was to create a "contingency" block; this will be discussed further in Section 2.6. For MGS, a partial solution was to set the data rate at the highest supportable rate in the weekly sequence but not to make tactical changes. However, this had unforeseen consequences, as detailed in Section 2.14.

*Lesson: Future orbiters should be designed to include sufficient flexibility to manage late requests by rovers or landers to reconfigure the relay link on a much*

*shorter timescale than that generally required for orbiter operations.*

### 2.2 Meetings

The Long-Range Relay Coordination (LRRC) Process was designed to overlay predictions for longer-term relay opportunities with predictions of the onboard activities of the relevant orbiters. Primarily this process was designed to gain an understanding of the overflight frequencies, estimated forward- and return-link data latencies, and restrictions by each spacecraft in the network; and to identify potential overflight conflicts during what became a six-week planning horizon for the relay activities.

The LRRC Meeting was held every four weeks ("monthly") and attended by all parties involved in relay coordination. The purpose of this meeting was to discuss the products and issues involved in the LRRC Process. This meeting was facilitated by automated conflict identification and resolution software. Typically, the meeting involved the discussion of the initial parameters used by the relay coordination software, and the (few) problems identified in the output products. This meeting never exceeded thirty minutes and often was completed in ten minutes.

Generally speaking, it was because the rovers landed at geometrically separated sites that there were few conflicts to resolve in the planning timelines. This, in addition to the quality of the software, caused this meeting to be less valuable than anticipated. Although the Long-Range planning products continued to be useful throughout relay operations, the meeting itself proved to be most useful during the first month or two of relay operations; after that, relevant Long-Range issues were discussed outside of this meeting or in the Short-Range Relay Coordination Meeting, as detailed next.

The Short-Range Relay Coordination (SRRC) Process was the mechanism to determine the details of specific relay opportunities. It allowed for a revision of the estimates for forward- and return-link data latencies based upon improved navigation predictions and an updated understanding of the orbiters' onboard activities. The rover operators used the results of the LRRC Process to guide them in selecting the overflights during which the rovers would request relay support. By doing so they were able to minimize or eliminate actual conflicts that occurred in the network.

This Short-Range Relay Coordination Process generated the "overflight plan" that represented the agreed upon actions that each spacecraft in the network

needed to perform to facilitate the relay activities. It occurred on a weekly basis to accommodate the changing needs of the landed assets while still retaining some form of long-term commanding that was required by the orbiting spacecraft.

A Short-Range Planning Meeting was held weekly with all parties involved in relay coordination. The purpose of this meeting was to discuss the products and issues involved in the Short-Range Coordination Process. The meeting never exceeded forty-five minutes and usually lasted thirty minutes or less. Typically, it involved the discussion of the initial parameters used by the software and (the few) problems identified in the output products. This process and associated software also worked well and, as a result, this meeting was valuable, but fairly routine.

*Lesson: The relay coordination processes and software tools should be emulated in future relay scenarios; the meetings were short and productive and the software served its purpose well.*

### 2.3 Infrastructure

In order to meaningfully support these meetings, it was necessary to develop an infrastructure that could meet the needs of each spacecraft's operations teams. This was particularly challenging since the coordination process needed to accommodate both NASA entities and the European partners who would be supporting the relay coordination task remotely, some of whom were functioning on "Mars time"<sup>4</sup> or in different time zones.

This infrastructure can be considered in three parts:

- 1) A consistent physical meeting location and teleconference numbers.
- 2) Properly functioning equipment.
- 3) Stable meeting schedules.

Because software plays an essential role in the relay coordination effort, it was necessary that a consistent meeting facility be established where personnel could become familiar with how the resident computer systems and displays operate. At JPL, it was often the case that different meeting locations had different types of computer equipment which made it difficult or

<sup>4</sup> Mars rotates slightly slower than Earth, making its day roughly twenty-four hours and forty minutes long. Operations teams for short-lived missions on the surface of Mars have historically operated the vehicles on cycles that match the Martian day in order to more efficiently utilize the lifetime of the vehicles.

impossible to display relay coordination products to the participants.

Properly functioning equipment for the relay operations meeting was also imperative. With regards to the computer system, it was occasionally necessary to rebuild relay operations products in real-time. In addition, the most effective means of communicating the overflight plan was to display certain graphics in the relay meetings. Without properly functioning computer equipment, this would not have been possible.

Functioning communications lines were also essential, especially static-free and reliable speakerphone capabilities to facilitate off-site attendees. This experience demonstrated the value of having a stable phone number for conference calls so valuable meeting time was not wasted in assuring all required attendees were present.

It was also important to establish a consistent meeting schedule. Inconsistencies in the meeting schedules could have resulted in confusion and schedule conflicts for key members of the relay operations team, particularly for those teams planning on "Mars time" or working in different time zones.

*Lesson: Consistent meeting times, locations, and teleconference numbers are essential. In meeting facilities, advanced attention should be paid to ensuring that meeting equipment is functioning properly and of good quality.*

### 2.4 Procedure Flexibility

Team members were guided by procedures to perform the relay coordination function, to interface with all the projects, and to take ownership of each "overflight plan." These procedures specified how the software should be operated, from whom to gather inputs, and to whom to distribute outputs. They also included information about how to run the relay coordination meetings. While these procedures were very thorough in their inception, it was discovered very early that they needed to be "living" procedures, with the flexibility to continually improve the relay coordination process.

These procedures were originally written prior to relay operations. Some steps of the procedures were captured from [2], a formal document that outlines the relay coordination processes. Other steps were added as a result of the test and training activities that occurred prior to relay operations. A single person was identified to be the owner of the procedures, and, as they were exercised by various members of the relay

coordination team, some steps would be suggested for inclusion or deletion. The owner of the procedure kept the official version and ensured that all accepted edits were passed on to the entire team.

*Lesson: Relay coordination procedures should be flexible and robust. Maintaining "living" procedures to adapt to the changing conditions of the network facilitates process improvement.*

## 2.5 Coordination

The software tools that were used to perform the coordination effort were originally designed nearly a year before relay activities commenced. These were built using the same strategy as in [4]. Therefore, many assumptions had to be made regarding how the network would function. For example, it was assumed that MGS would communicate at only one data rate, that the rovers would be relatively short-lived, and that the network would include all six planned spacecraft. All of these assumptions were eventually proved wrong, and it became necessary to modify the software to account for these differences.

Instead of being required to make these modifications at a late date, it would have been better had the software designers taken a broader perspective to account for more possibilities in the final configuration of the relay network. For example, the software was designed to automatically schedule support for every landed asset on every overflight except for those that were considered to be "too short," meaning shorter than a configurable duration that was an input parameter to the software.

In retrospect, it would have been beneficial to have had additional options that would have provided the ability to more selectively choose which overflights the orbiters supported. In this case, it would have been helpful if the software had been able to schedule support by an orbiter for only specified landers, as opposed to all of them; or to schedule support for an overflight that only occurred when the orbiter passed above the lander at a particular time of the lander's day. While the software did have several "knobs and dials" and performed well (particularly for a first-use situation), it was learned that more were needed.

*Lesson: Future designers of relay coordination software should engineer it to account for most possibilities in the network, utilizing past experience to guide them in determining meaningful and even least-likely options that may be utilized in the planning processes.*

## 2.6 Requirements Definition

A perpetual problem with software and procedure development efforts is the lack of well-established requirements. The challenge in this case was exaggerated because the effort involved the coordination of multiple projects which are individually considered operational and complete.

Operations personnel, in addition to their ongoing responsibilities, were involved in the development of the requirements for the relay activities. However, this requirements generation effort naturally needed to be of lower priority than ongoing operations activities. Even so, it was very valuable to involve these operations personnel in developing the requirements for the relay effort, since they were the ones most familiar with each spacecraft's operational complexity. Nevertheless, with these individuals frequently interrupted by daily operational concerns, and owing to the limited personnel and time horizon for relay development, it was virtually guaranteed that the initial requirements set would be incomplete.

This would not have been as much of a problem had the requirements definition effort taken place earlier and more intensively. Furthermore, while the requirements identified at the outset of the effort provided basic functionality, few, if any, of the people involved truly understood how successful the relay network would be. Only as the relay links proved to be spectacularly successful did the imaginations of the operations teams begin to realize the magnitude of what could yet be accomplished. This caused a post-development boost in creativity that resulted in a flood of additional requirements on the already existing system.

While many of these new requirements were minor, such as changes to planning data reports and formats, and adjustments to procedures; again the limited personnel available for development and ongoing operations work resulted in many small and unorganized changes to the system, rather than a coordinated and methodical approach.

One prime example of evolving requirements was the development by LMA of an alternative onboard command mechanism on Odyssey to perform relay activities, commonly referred to as the "contingency block."<sup>5</sup> This was in addition to the original

---

<sup>5</sup> The "contingency block" became the standard mechanism for performing relay operations on Odyssey. The production of new software to build and check sequences that used the block, as well as a new

commanding concept that had been designed, developed, thoroughly tested, and was ready for full implementation. As the Spirit rover approached Mars, the rover operators expressed doubt that an adequate prediction of the quality of the relay link could be determined in the longer-term planning timeframe required by the orbiting assets, and subsequently demanded more flexibility in the operation of the link.

This doubt led to the decision to use the alternative and initially unvalidated command mechanism on Odyssey which utilized specific settings of parameters, called global variables, to determine the behavior of an individual relay link. This capability is described further in [3]. This new strategy provided the ability to supercede a previously designed command sequence on a very short timescale. This provided a near-real-time override capability to control various aspects of the UHF link, such as the communication bit rate (the most common variable adjusted), the collection of 1- or 2-way Doppler data, forward-link command capability, and coding and modulation varieties in the event of off-nominal link performance.

This alternative command mechanism was used as the primary means of operating the relay link.<sup>6</sup> This eventually forced changes to the ground command development and modeling software, as well as the development of new procedures to manage this alternative strategy. Had more attention been given to the problem at the beginning of the requirements definition stage, the rover operators' need to command in near-realtime would have been identified early and the alternative capability likely would have become the baseline from the beginning.

As another example, MGS was originally intended to support only two overflights per day per rover. However, a new, late requirement arose that drove MGS to support every possible UHF opportunity during the Impact-to-Egress (ITE) phase. This late requirement was very controversial because it required the Mars Orbiter Camera (MOC) science team to collect relay data from the Mars Relay (MR) during all possible passes per day (up to a geometric maximum of ten per day) and to have buffer space available regardless of whether or not the rover team actually utilized the pass, resulting in a significant loss of MGS science data. In addition, the relay planning software had to be changed significantly to add flags, features,

---

round of testing, had to be completed in approximately two months, a very short time for a capability of this magnitude and complexity.

<sup>6</sup> The significant operational flexibility provided by the "contingency block" was rarely used in operations.

and commands only months before Spirit's Entry, Descent, and Landing (EDL) with very little time for testing. Also, the MGS relay sequences, which largely consisted of MR mode commands, became more complex.

The realization of these new and late-breaking MER requirements coupled with the intense NASA focus on MER mission success resulted in additional high-level management intervention. Progress slowed to service these new requirements amid the additional oversight. All of these late changes increased the overall work load on both the Odyssey and MGS teams supporting relay and ultimately increased operations complexity for the entire relay support period. Had the requirements been better understood more than just months prior to actual relay operations, the teams would have been in a better position to accommodate them.

*Lesson: Proper attention to requirements definition and operational scenarios early in development can prevent last minute and potentially costly and risky changes to the relay commanding paradigm. This should include a clear description of the expectations of the lander teams with respect to how they will intend to utilize relay services provided by the orbiter teams.*

## 2.7 Project Documentation

Each spacecraft team had their own unique methods of operating their spacecraft. One of the challenges of developing the relay capability was to come to an understanding of what the real requirements were and to make firm agreements between the projects. To this end, prior to relay operations more emphasis should have been placed on identifying and documenting inter-team requirements and agreements. Significant effort went into developing project-to-project Memorandums Of Understanding (MOUs), Interface Control Documents (ICDs), and Operational Interface Agreements (OIAs) which defined the details of the relay interfaces and responsibilities; nevertheless, as previously discussed, many details were not worked out until late in the development process and it was very challenging to keep appropriate records of all agreements that had been made and to keep the documentation up-to-date. In the case of Beagle 2 and Mars Express, this formal documentation was never signed by the authorities on those projects.

*Lesson: Future relay efforts should put more emphasis on the early generation and maintenance of the documentation which describes the inter-project agreements in order to account for and to minimize the cost of "creeping" requirements.*

## 2.8 Test and Training

Test and training activities associated with the relay process proved to be fairly unstructured because the individual projects drove the exercises. In the future, a level of control higher than any individual project would be useful to improve some of these activities.

Nevertheless, not all tests should be “graduated” to the control of a multimission entity. For example, thread tests<sup>7</sup> worked well with direct coordination between the individual projects. In the case of the Odyssey-to-Beagle 2 link, Odyssey drove the end-to-end thread tests and succeeded admirably in flowing data through the engineering models of the spacecraft’s onboard flight software, through a simulation of the transmission path from the spacecraft to Earth, and finally through Odyssey’s ground data system to the point where the data was moved to a server owned by the Beagle 2 team. Future projects would do well to emulate the Odyssey example in this case.

The Operational Readiness Tests<sup>8</sup> (ORTs) driven by the MER project serve as a counter-example. The MER EDL and Surface ORTs began late in the spring of 2003. The first few ORTs were internal to the MER project and did not include the relay component; this was necessary because the project used these tests to train the majority of the MER team members and to develop their detailed processes, which had not yet been defined. The first ORT to include the relay component occurred at the end of July 2003, and it was during this test that weaknesses in the relay process began to emerge.

The Mission Management Office’s (MMO) Mission Planning and Sequencing Team (MPST), tasked with planning and sequencing the activities of Odyssey and MGS, essentially led the planning associated with this first MER ORT that involved relay. In subsequent ORTs, the MER Mission Planning Team took the lead. In both cases, having personnel from one project direct another was an unproductive strategy, giving no one the authority to effectively manage the cross-project requirements and goals of the training exercises.

---

<sup>7</sup> Thread tests are designed to exercise the movement of data through all nodes in a project, usually in an end-to-end sense.

<sup>8</sup> Operational Readiness Tests are generally designed to exercise the processes that a project will utilize to operate a spacecraft. Particular emphasis is generally given to interfaces between teams and meeting schedules. These often serve as training exercises.

To compound the problem, as mentioned in Section 2.1, there were great mismatches between the planning horizons of the different projects. Both MGS and ODY already had standard processes in place for approving and sending commands and sequences to their spacecraft. The MER project operated under two planning schedules, a longer-term (1 to 2 weeks) “strategic” planning process and a shorter-term (1 to 2 days) “tactical” planning process.

For the later ORTs, the test coordinators on the MER project were “tactically-oriented,” meaning that their focus was on the day-to-day operations of the spacecraft as opposed to the advanced planning function. This caused difficulties when integrating the relay coordination processes, which clearly had a longer planning horizon than MER’s tactical processes. It was quickly shown that tactical personnel needed to be trained and represented in the strategic processes to ensure that required changes to the strategic plans could be facilitated in the simplest and most time-effective manner.

These experiences demonstrated the need for a test coordinator who could supervise the processes through completion. A successful test coordinator should have a high-level understanding of the orbiter uplink and downlink processes as well as the shorter-term commanding processes of the lander missions. In addition, this person would also make sure that each project was meeting their ORT objectives.

Yet another complication encountered during the pre-operations time period was that significant support for the extensive MER training activities was requested of Odyssey and MGS. However, since this work was not identified and scheduled in advance, it caused a large impact to the orbiter operation teams’ ongoing work, and in some cases support was simply not provided. Having this work identified, agreed to, and scheduled early could have allowed the orbiter projects to add additional workforce to cover the added activities.

*Lesson: It is recommended that the Mars Exploration Program office oversee these test and training activities via a designated multimission test and training director. This test and training director needs to have the high-level “big picture” of the entire process from the start of the orbiter planning to the downlink of the rover data, the authority to drive the process, and the responsibility to ensure that it is completed. This change should eliminate most of the problems observed in the test and training process.*

*Lesson: Agreements between the projects to exercise the relevant processes and to provide the appropriate*

*level of support for test and training activities should be established early in the development process.*

## **2.9 Inconsistent Management**

In this network there were two participants, Odyssey and MGS, which illustrated how all spacecraft operating in the relay network must be fully committed to participating.

The Odyssey project management was fully committed to participating in the relay network and to providing a reliable service to the landers. They actively participated in the coordination meetings which led to formalized inter-project agreements, stood behind those agreements, and worked to ensure that the relay service was supplied as agreed. Issues and concerns were clearly communicated to all participants in the relay network and care was taken to ensure that everyone understood the rationale behind any decisions that were made.

For MGS, the technical implementation of having the MGS relay data captured within the buffer of one of its primary science instruments had a direct impact on the science return of that instrument. This issue is discussed further in Section 2.14. Nevertheless, the complex interactions of the science instrument with the relay service were not the only factors that caused difficulties.

When the MER project realized they needed much more flexibility and support from the orbiters, a conflict between the MGS and MER projects occurred. The new desires of the MER project decreased MGS's science return, and since MGS was not designed as a relay asset like Odyssey, it was more difficult for the team to accommodate. Nevertheless, MGS management could have made some critical decisions early during this period on who had priority, MGS science or MER relay, which would have unified the MGS team to work through the problem and not work against each other. Alternatively, the Mars Exploration Program did not provide sufficient direction to clarify the role of MGS in the relay network.

What occurred was that the MGS project management tended to be less open and consistent with their decision-making, unwilling to commit to concrete requirements; and wanted to leave as much as possible open for later negotiation. This method succeeded in preserving as much science flexibility as possible, but caused endless problems for the relay development and operations team.

To highlight this, the original plan was to continue to operate MGS in the normal manner, with no special planning or commanding required. In that plan, it would have been up to the MER project to evaluate which MGS overflights to utilize based on the MGS plans for returning the relay data. MER was allocated up to two overflights per lander per day. This plan was originally agreed to by all participants.

However, late in the development and shortly before operations were to start, a new requirement was levied on behalf of MER to receive their data as quickly as possible, especially during their ITE phase. This resulted in a more complex sequencing strategy for MGS, which returned MER data to Earth more quickly, and more efficiently used the MGS onboard memory. Both of these results were good, but to enable them, major perturbations were required to the MGS sequence generation process as well as the relay plan. With this increased complexity, a week or more of planning and sequence generation was added to the normal timeline to develop the MGS command products. All these changes had to be completed in less than three months, a very short time for a capability of this magnitude.

A general lack in communication of the MGS project management's intentions, requirements, and decisions added confusion, delay, and rework to those performing relay development and operations. This caused a significant additional workload and stress on the relay team.

*Lesson: The project management for all spacecraft operating in the relay network must be fully committed to participating, be completely open and consistent with any issues or conflicts that arise, and adhere in good faith to the agreements that have been made.*

## **2.10 ITAR**

With the inclusion of the Mars Express Orbiter and Beagle 2 missions in the relay network, the United States' International Traffic in Arms Regulations<sup>9</sup> (ITAR) came into play. These regulations categorize spacecraft, specifically including scientific satellites, as defense articles subject to these controls. Any discussion of technical information regarding these spacecraft and their subsystems must be explicitly approved by the U.S. Department of Defense.

---

<sup>9</sup> The United States' International Traffic in Arms Regulations (ITAR) is a subsection of U.S. Executive Order 11958 - Administration of Arms Export Controls.

While JPL has some history of working with foreign technical partners, particularly with the remote stations of the Deep Space Network (DSN), major industry partners to JPL, such as LMA which helps to operate both Odyssey and MGS, are not automatically covered by the agreements that JPL has previously cleared through the State Department. Additional and individual agreements had to be formalized between each set of coordinating parties. Once the agreements were completed, all individuals involved in any technical interchange had to be properly briefed and/or trained in the extent of the technical data that could be passed. Penalties for violation of these regulations could have been severe, and could have resulted in the loss of contracts and significant monetary and criminal penalties for the parties involved.

The scope of these agreements, and the realization of the time needed to draft and approve them, was not thoroughly factored into the development schedules. Following the apparent loss of the Beagle 2 lander, a flurry of diagnostic activities occurred. During JPL-led teleconferences with the European personnel, LMA technical experts were prevented from participating in the discussions in any way because the necessary documents had not been approved or, in some cases, even written. Concern about the ITAR-allowed information prevented a full discussion of the technical details of the UHF radio link, such as the received power levels and other detailed performance data about the orbiting asset telecommunications equipment.

*Lesson: ITAR agreements should be written and approved long before relay activities ensue.*

### 2.11 Flexible Radio Design

The Cincinnati Electronics C/TT 505 radios used on Odyssey, MER, and Beagle 2 were designed with the settings shown in Table 1. These many options provided a high-level of flexibility in how the communication links between the spacecraft could be configured. This was especially useful for the Beagle 2

mission, where nothing was heard from the lander after its separation from the Mars Express Orbiter. Here, the Odyssey project was able to attempt a variety of configurations to eliminate possible Beagle 2 radio fault scenarios.

This flexibility was also used to good effect by the MER project in that Doppler, commanding, and data rates were all independent. This greatly eased the time spent scheduling and optimizing the configuration of the radios for each UHF overflight.

*Lesson: Future Mars relay spacecraft should include radios with sufficient communications options to provide a robust relay system. Too many options, however, can complicate the design and testing of the network, and increase operational costs.*

### 2.12 Flexible Sequencing Design

For the Odyssey spacecraft, the sequence that controls UHF communications sessions was initially intended to be independent of any other sequences that may be active. However, this approach was not feasible because Odyssey is unable to forward data to a lander at the same time that it is transmitting data to Earth. Because of this limitation, the sequence that controls relay activities must be overlaid and checked against the other onboard sequences to ensure that the relay activities do not conflict with other telecommunications and basic housekeeping functions. Validating that the relay sequence does so correctly takes additional time by Odyssey's ground operators. Were the relay activities truly independent of the nominal operations, less time and effort would be required to verify the safety and integrity of the relay command sequence.

In addition, there are other impacts to the network caused by halting the transmission of data to Earth to facilitate relay activities. The first impact is an increase in the return-link data latency for returning data from the landers via Odyssey for the duration of

**Table 1: Radio Communications Options**

	Odyssey	MER	Beagle 2
<b>Downlink Rates (Kbps)</b>	8, 32, 128, and 256	8, 32, 128, and 256	8, 32, and 128
<b>Command Rates (Kbps)</b>	8 and 32	8	2 and 8
<b>Receive Modes</b>	Reliable (Handshaking On and Off), Tone Beacon, and Canister		
<b>Receive Modulations</b>	Phase Shift Keying (PSK) and Frequency Shift Keying (FSK)		
<b>Receive Encodings</b>	Convolutional, Viterbi, Scramble, and No Coding		

the overflight. Second, the lack of a return-link means that lander data must be retained in Odyssey's onboard memory (which is limited in size) longer than anticipated such that a single overflight at the maximum data rate may overflow the available data volume. This will be discussed further in Section 2.20.

*Lesson: A relay system that functions independently from other spacecraft engineering and science activities can reduce the operational complexity of performing relay activities. This functionality should be incorporated in future relay spacecraft.*

### 2.13 Return-Link Data Prioritization

Another aspect that emerged as the Odyssey system was being optimized for high-volume data return was the need to be able to prioritize landed asset data for return to earth. The volume of Odyssey's onboard memory made available for each of the rovers was relatively small. To improve the use of that volume, it was decided to consolidate what were two separate buffers, one for each rover, into one single buffer. However, the Odyssey flight software design caused this buffer space to be transmitted to Earth in a first-in-first-out (FIFO) manner, without respect to the priority of the data. With high data volumes being received by Odyssey from the rovers, and decreasing data rates from the orbiters to the Earth (due to the increasing Mars-Earth range at the time), the data latency onboard the orbiter became a significant factor in the short-term planning cycle of the rovers.

*Lesson: Some method of ensuring the most important data is expedited and/or protected while being flowed through the relay network should be developed. This should be considered at a programmatic level as well, with the ability to prioritize data both between missions and within a single mission.*

Currently, there is no convenient way for the orbiting assets to protect high-priority relay data in the event of an Earth-communications failure. The onus is on the landed asset to protect the data until Earth receipt has been verified. A programmatic view should be taken as to whether this is the appropriate allocation of responsibility.

### 2.14 Radio as Facility Instrument

The MGS spacecraft design was inherited from the Mars Observer (MO) spacecraft, which had an MR UHF antenna. Contact with the Mars Observer was lost in August of 1993. After NASA gave the approval to build MGS, budget cuts occurred and it was decided that the MR, a spare from MO, would need to be

removed from the MGS payload to reduce hardware costs. Malin Space Science Systems (MSSS), the operators of the MOC on MGS, offered to attach the MR to the MOC data system, which enabled a simple and cheap method for MGS to retain the MR as a payload.

This enabled MGS to support future landed missions, but it unintentionally made the spacecraft more complex to operate as a relay asset. It was later learned that using the MOC instrument buffer for UHF relay data required detailed, accurate coordination between the MOC team and the MGS spacecraft, DSN scheduling, and sequencing teams. Data received or generated by the MR were stored within the MOC buffer and returned to Earth as MOC data. On Earth, the MOC ground data system was responsible for extracting the data acquired by the MR from the MOC data stream, processing the MR housekeeping data, and providing it back to JPL in a timely fashion.

In addition to this, significant time and planning was devoted to determining the amount of space in the MOC buffer that the MOC team was willing to provide to the MERs. Every bit provided to MER was one less bit available for MOC science. This was a source of contention between the two projects and the situation was exacerbated because neither team, MER or MGS, had time to discuss and understand the science impacts if MER was allowed to use the MOC buffer to the maximum extent. As stated in Section 2.9, MGS or Mars Exploration Program management could have helped this situation by prioritizing the MGS mission priorities during MER prime mission; instead, a lot of time was spent discussing possible compromises and analyzing individual scenarios.

Ultimately, a compromise was reached on the amount of buffer usage, but in order to minimize the impact on MOC science, reduce return-link data latency, and maximize the data allocation available for MER, the MGS spacecraft and sequencing teams were required to customize the orbiter sequences. Also, the times at which MGS would be scheduled to downlink data to the DSN had to be increased and carefully timed to ensure the proper flow of data. This forced the creation of a new strategic planning process that included forecasting DSN tracking time, which had to occur before the orbiter's sequences could even be built. This was an iterative process of manually adding, deleting, and trading DSN coverage as well as defining, on an orbit-by-orbit basis, what data type MGS would be transmitting to Earth.

Also, the MOC buffer was commanded exclusively by the MOC science team, giving them the ability to

decide how the data in the buffer was transmitted. The MOC team had two options for returning data to Earth: Science and Engineering 1 (S&E1) and Science and Engineering 2 (S&E2). In the S&E1 mode, the data was sent to solid state recorders (SSR) with long return-link data latencies. In the S&E2 mode, the data was transmitted to Earth immediately – a faster, but less dependable route. Simply clearing the MOC buffer could take thirty to forty-five minutes for the S&E2 return versus three or nine hours for S&E1 data to be returned at MGS's high or medium X-band data rate, respectively.

The MMO MPST used the software tools mentioned in Section 2.5, as in [4], to automatically determine the orbiter return-link data latencies based on the MGS background sequence. However, the calculations performed by the software were not able to determine which link, S&E1 or S&E2, the MOC team would choose to return the UHF data to Earth. The tool was unable to automatically determine how the data would be returned since there were many different interactions going on in the MOC buffer of which only the MOC science team was aware. In addition, the MGS science planning was done on a later timeline, well after the weekly relay planning process was concluded. Also, many of the S&E1 calculations for MGS were very rough because it was difficult to say where the UHF data was located within the MOC buffer. Those details were necessary to calculate high-fidelity return-link data latency times.

*Lesson: Having the UHF data flow through a science instrument buffer significantly increased the complexity in planning and implementation, as well as caused conflict between the two users. Future projects would do well to avoid this situation.*

## 2.15 Science Representation

Due to the fact that MGS was forced to transfer relay data through the MOC buffer, the process of acquiring orbiter images and performing relay passes via MGS became inseparable. While this was the source of much difficulty in actually commanding relay and science activities, it proved advantageous for coordinating science return.

MOC science team members, in addition to both strategic and tactical representatives from MER, attended the relay coordination and planning meetings. In this manner, coordinated planning occurred for UHF overflights, resulting in the maximum number of relay passes and overflight MOC images being obtained.

*Lesson: Having science representatives from both the landed asset and the orbiting asset attend the relay coordination meetings could provide a mechanism for improving overall science return.*

## 2.16 Radio Quirks

The MGS MR can only transfer data on the return-link (surface asset to the orbiter), but the forward-link is necessary to operate the link in the MBR (Mars Balloon Relay) protocol mode. The MBR protocol was designed to relay telemetry received from MBR-compatible Mars surface asset(s) back to Earth. The protocol is based on the use of a calling sequence, named Balloon Telemetry Time Slot (BTTS), which is 16 seconds long and is continually transmitted by the MR. Once a surface asset responds, the MR switches modes and starts to output its housekeeping telemetry (HKTM) and Doppler data. This means that the MGS UHF radio uses about one second of every sixteen during transmissions to insert Doppler and HKTM data. In the case of MER, the rovers continuously transmitted data, ignoring the MBR protocol. While this allowed continuous MER transmissions by eliminating the need for handshaking, it caused the MR to "lose" approximately one second of MER data out of each sixteen seconds transmitted. Any data received by MGS from MER during these gaps was lost.

This loss of data was expected and known on both the MGS and MER projects. However, the MER project did not realize the full impact of the periodic outage. Once surface operations began, the MER team had to deal with the added workload of retransmitting the periodic data losses due to the MBR protocol, which they had not planned for. In addition, many of the MER science teams were not fully informed of the impacts and were surprised to see holes in their data. Some members on the MER team did not understand why these gaps in the MGS-returned data existed and incorrectly assumed that there was something wrong with the MGS communications links.

In this experience, there were several different radios with different protocols that were compelled to function together. In the future, it is likely that a similar situation will exist and understanding these differences and designing the overall system to accommodate them will streamline processes and aid in the return of "good" and complete relay data.

*Lesson: The Mars Exploration Program office should create and maintain a database of known idiosyncrasies of the participating hardware, as well as existing project procedural idiosyncrasies and/or limitations, which should be a required input to the*

*design and implementation of future relay-capable missions. This will allow new projects to accommodate these idiosyncrasies at the earliest level of their development, reducing late-adaptation cost, schedule, and procedural changes.*

## 2.17 DSN Coverage

The effectiveness of the relay link was greatly enhanced by having relatively stable and ample DSN station coverage for the orbiters. At first, this was not the case as the MER project continually requested changes to the DSN allocations, attempting to optimize their DSN coverage. This forced the relay teams, especially for the orbiters, to spend significant amounts of time updating the weekly plans and sequences and coordinating changes with all participants in the relay link. As the MER project learned that their requests for these changes often adversely impacted the other projects and put their own data at risk, requests to change the DSN schedules became infrequent. Eventually, it was learned that the near-continuous DSN coverage for Odyssey allowed for very short return-link data latencies, providing a stable and reliable alternate route (other than direct-to-Earth) for the MERs' tactical data. As MER gained confidence in the relay process, additional high-priority data was sent via the UHF link.

MGS also benefited by having more DSN coverage than what was originally required to obtain normal MGS playback data. This additional coverage allowed MGS to schedule real-time telemetry after most overflights, reducing the return-link data latency by avoiding storage of the relay data on the Solid State Recorders (SSRs), which would have increased the latency from approximately one hour to up to 38 hours.

One of the methods used for increasing downlink coverage for Odyssey and MGS was to employ the DSN Multiple Spacecraft Per Aperture (MSPA) capability. This allowed a single DSN antenna to track two spacecraft simultaneously; however, in this configuration only one spacecraft at a time can use the ground transmitter. Despite this, through careful planning and use of non-MPSA configured antennas, both Odyssey and MGS were provided with adequate uplink capability. This was particularly important for Odyssey because it was the primary relay asset for the MERs and provided enhanced flexibility in managing relay opportunities to increase data return while maintain short latencies. This flexibility worked well to complement the shorter MER planning cycle.

*Lesson: With the proper scheduling of DSN time to provide sufficient coverage to the orbiters, forward-*

*and return-link data latencies can be dramatically reduced. This results in more efficient data return for both the landed and the orbiting assets.*

## 2.18 Realtime Operations

Since the lander activities were tied to the Martian day, which is about forty minutes longer than an Earth day, and with three landers scattered across the face of Mars, key activities could happen at any time of the Earth day. Therefore, with the development of the Odyssey "contingency block," as discussed in Section 2.6, the decision was made to provide around-the-clock realtime operations support<sup>10</sup> from the landing of Beagle 2 through the MER ITE period. This persistent realtime operations coverage proved invaluable during the early Beagle 2 mission when the flexibility of the contingency block was put to the test by varying the relay parameters in the attempt to contact the lander.

In addition, the initial scheduling for the MER relay links included conservative choices for return-link data rates (32 kbps). The success of relay operations for MER was demonstrated early on, and confidence in the UHF link increased to the extent that the block parameters were exercised to increase the return-link data rate to Odyssey to 128 kbps on January 6, 2004; the second day after Spirit landed. The data rate via MGS was increased a short time later to the same rate. This four-fold increase in data volume provided the rover science and engineering teams with a flood of information on which to base their decision-making processes during the early phases of characterizing the rovers' systems in a new environment.

*Lesson: To provide the greatest robustness in a similar relay network, future projects should provide flexibility in the orbiter systems to allow near-realtime changes in state, request continuous ground station coverage for telemetry and commanding, and supply round-the-clock realtime operations support. These are especially important during the early commissioning and characterization phases of a lander mission.*

## 2.19 Forward-Link Capability

Odyssey faced a difficult challenge in the forward-link path; all three landers could have used Odyssey's forward-link capability on any overflight. A process

---

<sup>10</sup> Realtime operations support implies that there are personnel available to facilitate the transmission of commands to a spacecraft. In this example, the continuous support provided for realtime operations was provided by the LMA realtime operations team.

was required to allow forward-link files to reside on-board Odyssey that could be sent to the landers at the correct time and in the correct order, but transmitted to Odyssey at any time. In addition, this needed to be done in such a way as to reduce the impact to the Odyssey staffing and budget profiles.

These problems were resolved by treating the landers as Odyssey science users, allowing the files for forward-link commanding to be submitted exactly as the Odyssey science teams submit their instrument commands. Science commanding on Odyssey is accomplished using a proven Mission Management Office (MMO) capability called the Non-Interactive File Load (NIFL) Process, as in [5]. These files are transmitted to the spacecraft as binary data files, the contents of which the spacecraft ignores, and stored in the spacecraft's memory. The mechanism for performing the relay activities onboard Odyssey anticipates the presence of these forward-link command files, which must be properly named and resident in the spacecraft's memory prior to the overflight. This design was detailed in an earlier paper, referenced in [1].

This implementation was highly capable and affordable since it was a simple extension of an existing process within the Odyssey flight team. In practice, the system worked very well. Significant testing of the forward-link was performed with the MERs using large files in anticipation of potential lander flight software reloads performed through Odyssey. To date, the MERs used the forward-link operationally about 10 times. In all cases the ground system and flight system interfaces worked exactly as anticipated.

*Lesson: It is likely that orbiters will always be required to implement forward-link paths since they are viable contingency command paths. Future missions should emulate the Odyssey implementation in order to provide a robust system at a low cost.*

## 2.20 Onboard Storage

On both orbiters, the amount of onboard storage for the relay data was one of the key limitations for overall data return. Since Spirit and Opportunity landed roughly 180 degrees apart in latitude and with the orbiters in near-polar orbits, back-to-back overflights separated by an hour were a common occurrence. The lack of available memory on the orbiters increased operations complexity because two Mars rovers were now competing with each other for the same memory allocation. Additionally, orbiter science was simultaneously being collected, which also required memory.

For Odyssey, additional onboard storage space would have eased operations. Although Odyssey far exceeded the minimum required data return for the MER missions, it was found that the available storage space in Odyssey's packet buffer could be filled up in roughly two "good" overflights.

Pre-flight, the lower 128 kbps return-link data rate was thought to be the maximum reasonable rate and the packet buffer was sized to accommodate one overflight at this data rate. The radios performed better than some expected and a rate of 256 kbps was ultimately achieved. However, if there was a sufficiently long gap in DSN coverage, then there was no means for Odyssey to downlink the rover data and to clear out the packet buffer prior to another overflight. In practice, the problem only became noticeable on Odyssey when the highest 256 kbps rover-to-orbiter data rates were used, a situation exacerbated by Odyssey's increasing Mars-Earth range and subsequently lower data rates to Earth. The packet buffer limitation drove changes to the flight and ground systems to better optimize the available packet buffer space. The situation would have been more complicated if Beagle 2 had survived. In hindsight, it would have been valuable to have sized the packet buffer for the worst case loading of all three landed assets.

Along similar lines, the MGS spacecraft has four solid state recorders (SSR) that all the instruments share. Instead of allocating a certain amount of space for each instrument to use, MGS continuously draws data from all the instruments and combines it into one stream, which is either recorded (S&E1) or sent directly out the X-band link (S&E2), as previously discussed in Section 2.14. With the MR connected to the MOC buffer, one "decent" overflight at 128 kbps could easily fill over 85% of the buffer space. Even when MGS is transmitting the data to Earth at its medium data rate (the highest rates being unachievable due to the Earth-Mars range), about forty-eight minutes was required to clear out data from the MOC buffer. This was potentially problematic when several overflights were spaced only an hour or less apart because it required MGS to always have a DSN antenna available to receive the data.

To simplify and prevent buffer overflow, the MER team did not request back-to-back overflights with MGS. For example, they typically requested a pass for Spirit during its local night and then about twelve hours later they requested an Opportunity pass during its local night. In addition, since the buffer was a shared resource between the MOC and the UHF relay data, the MGS sequences were optimized to empty the buffer of

any UHF data as soon as possible so that the MOC team could continue to collect Mars mapping images. This optimization required DSN tracking on or soon after UHF overflights, new high-level strategic planning charts, as well as custom-built sequences for the spacecraft. If these conditions were not met, MGS science return was reduced to an unacceptable level.

*Lesson: Future orbiting assets could greatly simplify their onboard relay support by providing ample, designated, and partitionable onboard storage to facilitate both return- and forward-link products from multiple landed assets.*

### 3.0 Conclusions

The experience of coordinating relay activities between the diverse and complex spacecraft at Mars has been valuable and educational. While the lessons detailed in this paper are not exhaustive, several main themes can be derived from them, namely:

- 1) Relay systems onboard future spacecraft should contain flexible (but not too flexible) radio systems, be operationally separate from other onboard activities, and have sufficient data volume provided within them to facilitate both the return- and forward-link of any reasonable scenario. Similarly, ground data systems for each of the relay spacecraft should provide flexible architectures to facilitate the operation of the network and for the flow of data to-and-from other projects, both locally and remotely.
- 2) The true relay requirements for future Mars spacecraft should be defined early in their design phase, and these requirements should be exercised in ORTs and other testing and training activities.
- 3) The Mars Exploration Program office should designate visible and authoritative managers to coordinate both the relay requirements definition efforts across projects and to facilitate all test and training activities.

Many of the lessons outlined herein are already being incorporated into the design and operational plans of future Mars spacecraft and it is expected that by doing so the full potential of the relay network will continue to be realized.

### 3.1 References

[1] R. Gladden, P. Fieseler, B. Waggoner, M. Thornton, R. Thomas, and P. Hwang. "Coordinating UHF Relay Activities at Mars", Proceedings for the IAA Fifth International Conference on Low-Cost

Planetary Missions, Noordwijk, The Netherlands, September 2003.

[2] R. Gladden. JPL Document D-26272, Mission Management Office Relay Operations Plan, Version 1, 6 May 2003.

[3] C. A. Grasso, "The fully programmable spacecraft – Procedural sequencing for JPL deep space missions using VML (Virtual Machine Language)", 2002 IEEE Aerospace Conference Proceedings – Volume 1, Big Sky, MT, March 2002, Piscataway, NJ, IEEE, 202, p. 1-75 to 1-81.

[4] R. Gladden, "AUTOGEN: The Mars 2001 Odyssey and the 'Autogen' Process", 2002 AIAA/USU Small Satellite Conference Proceedings, SSC02-IV-2, Logan, UT, August 2002.

[5] R. N. Brooks Jr., "A High Efficiency System for Science Instrument Commanding for the Mars Global Surveyor Mission", RAL.GS.36, Proceedings of the First International Symposium on Reducing the Cost of Ground Systems and Operations, Rutherford Appleton Laboratory, Chilton, Oxfordshire, UK, September 1995.

*The work described in this paper was performed at the Jet Propulsion Laboratory (JPL), managed by The California Institute of Technology (CalTech), under contract to the National Aeronautics and Space Administration (NASA).*