IAC-04-Q.3.a.03

MARS EXPLORATION ROVER: LAUNCH, CRUISE, ENTRY, DESCENT, AND LANDING

J. K. Erickson[1], R. Manning[2], M. Adler[3]
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California, USA

ABSTRACT

The Mars Exploration Rover Project was an ambitious effort to land two highly capable rovers on Mars and concurrently explore the Martian surface for three months each. Launched in June and July of 2003, cruise operations were conducted through January 4, 2004 with the first landing, followed by the second landing on January 25. The prime mission for the second rover ended on April 27, 2004. This paper will provide an overview of the launch, cruise, and landing phases of the mission, including the engineering and science objectives and challenges involved in the selection and targeting of the landing sites, as well as the excitement and challenges of atmospheric entry, descent and landing execution.

1. INTRODUCTION

2003 will be remembered in history as the year mankind sent twin robotic explorers to the surface of Mars. The pair of solar powered six wheeled rovers, with their advanced remote sensing capabilities and ability to place in situ instruments onto rock and soil surfaces with their instrument arm, were going to revolutionize the scientific thinking about the history of water on Mars. Along the way, they would capture the imagination of the world, both by showing that NASA could recover from the Mars '98 failures with a bold risky mission and by allowing the public to see the human side of mission operations in "real-time".

The original plan for this exciting mission has previously been described[1]. The intent of this paper is to cover the time period between the start of launch preparations in May 2003 and the second landing in January 2004. Since the rover design, and the instrument complement they carry is described in the 2002 report to the IAF[1], it will also be excluded from this report. An additional paper to be presented at this same conference[2] will cover the surface operations phase.

2. MISSION OVERVIEW

The mission plans of the twin Mars Exploration Rovers (MER) were virtually identical, with the exceptions of the targets and the launch and arrival dates. June and July 2003 marked the successful launches of the vehicles, with both on track to their respective targets, Spirit to a landing site to be designated

[1]Mars Exploration Rover Project Manager, [2]Mars Exploration Rover EDL Manager, [3]Mars Exploration Rover Spirit Mission Manager

later, and Opportunity off to Meridiani Planum. As will be discussed later, the final targeting of Spirit to Gusev Crater was reserved until after the successful launch of Opportunity. The roughly seven month cruise period stretched from the successful launches, through pre-Entry, Descent, and Landing (EDL) preparations, to the actual EDL events. The EDL events themselves lasted only a few hours, beginning with the turn to entry attitude.

3. LAUNCH PHASE

Two different subjects are significant to the launch phase of MER. The planned launch period design, and the actual launch period utilization. These will be discussed in the following paragraphs.

3.1 Launch Period Design

MER was required to launch two rovers in the 2003 Earth-to-Mars opportunity on separate Delta 2 launch vehicles. The small amount of time in 2003 available for minimum-energy trajectories, logistical constraints on the launch of two Delta 2's, the required arrival conditions for the trajectories, the need to maximize the allocated spacecraft mass, and delays during the actual launch campaign all combined to make the launch strategy a difficult and dynamic problem.

The Earth-to-Mars launch opportunities in 2003 were well suited to the MER mission in several respects. They are among the lowest energy transfers in the 32-year cycle, they provided arrivals in the late Southern Spring allowing for a solar-powered mission in the Southern latitudes which contained highly desirable science targets, and they facilitated prograde, low atmosphere-relative velocity entries with Earth visibility of entry, descent, and landing for critical communications during that phase. A Boeing Delta 2 7925 launch vehicle was selected when the MER project was originally planned to be a launch of a single spacecraft. The launch period was selected to maximize the available mass for the spacecraft, to be at least 18 days long, to have a constant arrival date, and to provide at least a 10° Earth elevation five minutes after landing for entry, descent, and landing communication. A launch period of 18 days was required to provide a 99% probability of initiating a launch with respect to historical weather, range, and launch vehicle delays. Meeting all of those constraints resulted in a launch period from May 30th through June 16th, 2003, with a constant arrival date of January 4th, 2004. The 10° Earth elevation constraint determined the earliest constant arrival date. Given all that, the minimum energy determined the start and end dates, which resulted in a maximum C3 of 9.4 km^2/s^2. The trajectories all easily met a constraint on the arrival v-infinity of 3.0 km/s, as well as launch declination requirements for the Cape Canaveral launch site in Florida.

These characteristics are illustrated in the launch/arrival date plot below (Figure 1). The key contours shown are launch energy in black, Earth elevation five minutes after landing in blue, and arrival v-infinity in green. The constraint contours for the various parameters are in bold. The final launch periods are the thick red and green bars marked MER-A and MER-B. The original single-MER launch period is the first 18 days of the MER-A bar.

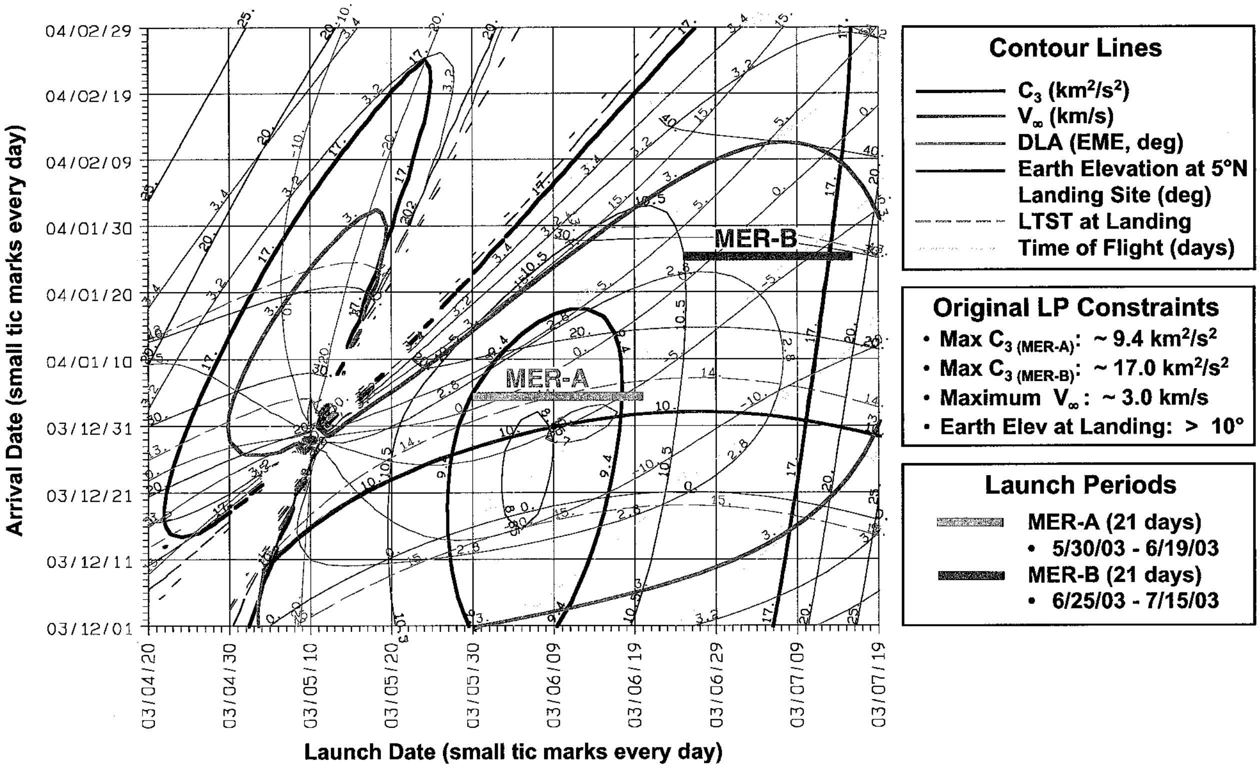

Figure 1. Final Launch Periods

When a second MER spacecraft was added, also to be launched in the 2003 opportunity, the first one remained on the Delta 2 7925. A more capable Delta 2, a 7925H with larger strap-on solid rockets than the 7925, was required to provide a comparable mass capability for a second launch later in the 2003 opportunity. An arrival date of three weeks after the first was selected to balance science return, which is reduced as the arrival moves later in the Martian year, against the ability to have enough time to react to problems with the first landing and critical deployments and possibly reprogram the second vehicle before entry. This put the constant arrival date of the second mission, designated MER-B on January 25th, 2004, 21 days after the arrival of MER-A on January 4th.

Due to logistical constraints on turning around Delta 2 launches, it would require at least 10 days after the actual launch of MER-A to prepare for the first launch attempt of MER-B. If the start of the MER-B launch period were less than 10 days after the end of the MER-A launch period, then a very late MER-A launch would prevent the use of the first day or few days of the MER-B launch period. As a result, it was desired to minimize that overlap. Since the two MER vehicles were designed to be identical, the other constraint that determined the MER-B launch period was to select a launch energy that provided the same mass capability on the 7925H as 9.4 km^2/s^2 would provide on the 7925. That turns out to be 17.0 km^2/s^2. The 18-day MER-B period was placed to end at the C3 = 17.0 contour, which slides it as far right as possible to minimize the logistical overlap between the MER-A and MER-B launch periods. This set the launch dates for MER-B to be June 25th through July 12th, 2003. The resulting logistical overlap was only two days. Since the probability of delaying that far into the MER-A period, combined with the probability of also delaying the MER-B launch enough that the loss of the first two days mattered was very small, this overlap and possible loss of the first one or two days of the MER-B period was acceptable.

Given these launch periods and launch vehicle capabilities, the spacecraft mass allocation was set to 1077 kg to provide a 99% probability of commanded shutdown (PCS), i.e. not running out of propellant on the Delta 2 second stage, for the optimum launch azimuth of 93° across the launch periods. A second launch azimuth of 99° was added for each day to provide a second chance about 40 minutes after the first in case weather cleared up or range problems were resolved in that time. Lower PCS performance was accepted for those azimuths on the open and close days of the MER-A period and the close day of the MER-B period, which were at the C3 limits. For the Delta 2, each launch opportunity at a given azimuth is instantaneous, so the second azimuth significantly opens up the window for that day.

The choice of 18 days for the launch period was made based on historical data of approximately 50 Delta 2 launches up to 1996. A more complete set of historical data was analyzed that included later missions bringing the total to 107 Delta 2 launches through 2001. This resulted in a new estimate of a 20-day launch period required to achieve a 99% probability of launching with respect to weather, range, and launch vehicle delays. The additional days were accommodated by adding launch days at the end of each period, and by accepting slightly lower PCS performance on those days in order to retain the existing mass allocation for the spacecraft. Three days were added to each, resulting in two 21-day launch periods, with two azimuths per day. The 21st day was added to each to have those target designs complete and validated in case a late release of launch vehicle performance margin made those days viable. Days were added to the end of the MER-B launch period instead of the beginning to provide as late a launch as possible for schedule resiliency. The logistical overlap was increased by two days for the required 20-day periods, but again this was acceptable due to the low probability of both missions experiencing extreme delays. The lower PCS performance of the appended days would be evaluated for each azimuth once the final launch vehicle performance was established. A minimum acceptable PCS of 95% was established for the contingency opportunities. The 93° launch azimuth on the first 20 days of each launch period was assessed during design to be greater than a 97% PCS, and so the two 20 days periods satisfied the project requirements.

During the development of the final launch vehicle targeting, options were developed for higher spacecraft mass allocations to provide additional margin if needed. Based on updated knowledge of the launch vehicle component masses and performance margins, the final targets used 1081.5 kg as the allocation for MER-B, with MER-A remaining at 1077 kg.

The final launch periods were May 30th through June 19th, 2003 for MER-A, and June 25th through July 15th for MER-B, with arrivals at Mars on January 4th and January 25th, 2004 respectively. Constraints on the spacecraft attitude with respect to the Sun and the spacecraft battery lifetime and Earth shadowing shortly after launch resulted in the selection of a daytime liftoff, short-coast orbit for all MER-A launch opportunities and a nighttime liftoff, long-coast orbit for all MER-B launch opportunities.

3.2 Launch Period Utilization

Part of the spacecraft mass allocation was the propellant used to correct for launch vehicle injection errors and to target the spacecraft to the desired atmospheric entry at Mars. The tanks however were large enough to hold 7 kg more propellant than was allocated. It was desirable to fill the tanks if possible to provide margin against unexpected spacecraft events or a launch vehicle un-commanded shutdown resulting in a limited C3 shortfall. During the development of the launch vehicles, a ballast allocation is maintained to account for uncertainties in the mass and performance of the launch vehicle components. This ballast is set conservatively to assure that the launch vehicle will meet the contracted performance requirements. In the event that the launch vehicle would require a significant amount of ballast, an arrangement was made to have the spacecraft provide some of that ballast in the form of spacecraft propellant. This would allow the tanks to be filled and increase the probability of a successful mission against unexpected events. In the end, the spacecraft itself came in light, and so the tanks were filled without having to substitute propellant for launch vehicle ballast. In fact, the launch vehicle ballast was increased to make up for the lighter spacecraft. There are limits to how much ballast can be mounted on the launch vehicle. This was taken into account in the planning for the mass targets using optimistic assumptions for final launch vehicle and spacecraft masses.

As the first launch day of MER-A approached, the spacecraft and launch vehicle were stacked and ready to go, but further testing, analysis, and reviews of the spacecraft design were required before authorization to launch. The first launch attempt was scheduled for June 8th. The spacecraft-related delay cost the first nine days of the launch period. During this time, a previously developed option was exercised to provide additional launch period days by planning to remove launch vehicle ballast in the event MER-A did not launch by June 19th. Approximately 19 kg of ballast was available for removal, which provided enough performance to add five days to the launch period, extending it to June 24th while maintaining better than a 95% PCS for the last day. The ballast could be removed in less than a day, allowing for a June 20th launch after a June 19th scrub. The new MER-A launch period was June 8th through June 24th.

The launch was authorized, and the teams initiated the countdown procedure for a June 8th launch. Thunderstorms in the vicinity prevented launch on the 8th and 9th of June. On the morning of June 10th, the weather was clear and MER-A, renamed "Spirit" a few days earlier, was launched on the first azimuth at 1:59 p.m. Eastern Daylight Time. The launch injection to the interplanetary trajectory was well within the expected uncertainty. The spacecraft performed a relatively small maneuver ten days later to correct the trajectory, using only 15 kg of the 52 kg propellant available. In fact, most of that 15 kg was used to correct a deliberate bias in the trajectory for planetary protection. Less than 5 kg of propellant were used for the remainder of the transit to Mars, mostly to remove a landing site selection bias.

The MER-B launch, scheduled for June 25th, was delayed for three days to allow time to complete the preparation of the launch vehicle, including repairs to the first stage insulation. The first launch attempt occurred on June 28th. A small boat on the range prevented a launch at the first azimuth. The weather got worse by the time of the second azimuth, and the launch was scrubbed due to high-altitude winds. An inspection of the first-stage insulation revealed that some of the insulation had de-bonded. Repairs as well as testing of the bonding technique and the time to cure the adhesive were required. While the repair and testing was ongoing, three more contingency days of launch targets were developed and validated, for July 16th, 17th, and 18th (UTC). The use of those targets would require the removal of 22 kg of ballast on the launch vehicle, later arrival dates at Mars by one, two, or three weeks respectively, and the acceptance of PCS values as low as 95% for targets with C3's up to 18.5 km^2/s^2. It was decided that the ballast removal would be performed for the July 14th and 15th launch dates as well, using newly developed targets in the event the launch were delayed past July 13th. This would significantly improve the PCS for those opportunities.

The insulation repairs delayed the launch until July 8th UTC (July 7th EDT). Before then however, a battery cell in the flight termination system showed degradation and required replacement. The battery was replaced, adding another 24-hour delay and taking the first launch attempt to July 9th UTC (July 8th EDT). In total, there were 12 days of launch vehicle delays and one day of range/weather delays. Finally on the evening of July 8th (EDT) the weather was good, and the teams went through the countdown procedure for the first azimuth, which proceeded as expected until eight seconds before launch. At that time an oxygen valve failed to close when commanded, and the launch was aborted. The problem was subsequently resolved, and the vehicle was recycled for launch on the second azimuth. MER-B, named "Opportunity" successfully launched at 11:18 p.m. EDT on July 8th. The MER-B launch injection performance was very good, requiring from the 52 kg available only 12 kg of propellant to correct the trajectory ten days later, including the correction of the planetary protection bias. On the order of 1 kg of propellant was used for the remainder of the transit to Mars.

This launch campaign ably demonstrated the value of both long launch periods and more than one launch opportunity a day. In both cases, the launch took place towards the middle or end of the available days, and for MER-B, the second azimuth was utilized. The experience also demonstrated the importance of being prepared to take advantage of late knowledge of performance margins. For both launch periods, long delays prompted the development of contingency extensions to the launch period that took advantage of ballast removal while on the pad, as well as later arrival dates for MER-B. Careful assessment of the resulting PCS for all options guided the launch period design and decisions on late extensions.

4. CRUISE PERIOD

The following sections will discuss the period after the launch, but prior to the Entry, Descent, and Landing (EDL) events. The project used the acquisition of signal by the Deep Space Network antennas immediately after launch as the boundary definition between launch and cruise. There was not a crisp division between the events described as cruise versus EDL, but the turn to entry attitude can serve as a reference for further discussion.

4.1 Cruise Operations Concept

The operations concept for the Mars Exploration rover Project called for a disproportionate effort to be spent on planning and training for the surface preparation phase, as opposed to cruise operations. The goal was to spend time and resources addressing the more unknown and riskier portion of mission operations, with less emphasis on preparing for a "quiet cruise". In general, this approach proved worthwhile. Key exceptions were the early emphasis on the navigation of the two missions. Navigation had been singled out early as a challenge, both due to concerns about a previous problem with the Mars Climate Orbiter, and due to the stringent accuracy requirements driving the landing site selection choices. As discussed later, the Navigation plan for this mission relied on the delta Differenced One-way Ranging (dDOR). This data type proved to be mission enabling – allowing us to certify the Gusev Crater landing. Without dDOR, the landing dispersions would have been so large as to prohibit the selection of Gusev.

At the mission planning level of abstraction, the cruise period for each rover was broken into 6 high-level periods. Each period contained many individual activities, only the most major of which will be discussed in the succeeding paragraphs.

4.2 Cruise Timeline

Figure 2 shows a timeline of events for the cruise period, including depicting the final launch periods, through cruise, approach and finally ending with EDL. The major activities shown on the timeline will be discussed below, specifically including TCMs. Additional discussion follows on the various types of anomalies that occurred during cruise.

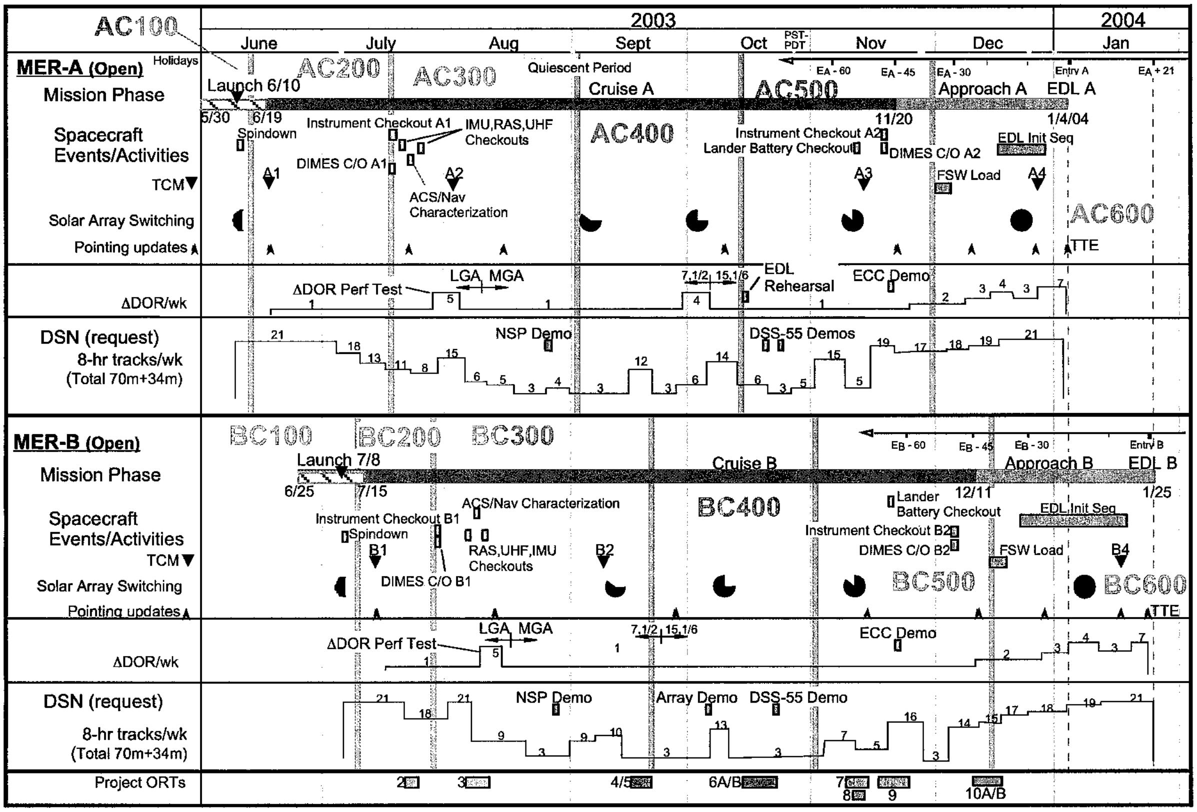

Notes: 1) Not shown are anomaly-related and regularly recurring activities, including uplinks, monthly HRS maintenance, and ~weekly Mini-TES Xcal and Delta-DOR
2) DSN coverage is shown as planned; some late changes are known to have occurred
3) All dates UTC-SCET

JL
20 Aug 04

Figure 2. Cruise Activity Timeline

4.2.1 Trajectory Correction Maneuvers

Trajectory Correction Maneuvers (TCMs) were planned at intervals during the mission so as to meet two main requirements: the need to postpone the final decision on which landing site to target Spirit to until after the successful launch of Opportunity (discussed further in section 5.3) and, the need to achieve and maintain the necessary navigation accuracy to keep the rovers on a safe atmospheric entry path and one that would have the vehicles land at the designated landing site. As a part of the Mission Design process, there were seven windows identified to execute maneuvers. Four were normal windows, but TCMs 5 and 5X were designed such that only one of the two would occur. If TCM 5 were to be executed, then TCM 5X would not. Similarly, if TCM 5 could not occur (due to a ground or spacecraft problem), then TCM 5X would be executed. TCM 6, on the other hand, was to be used only in the event of discovering a large orbit determination error on final approach to Mars.

TCM	Date (UTC)	Total Delta V (m/s)	Magnitude Error
A1	6/20/03	16.460	-0.70%
A2	8/1/03	6.008	-2.40%
A3	11/14/03	0.577	-1.70%
A4	12/27/03	0.025	-1.40%
A5	1/2/04	Canceled	
A5X	1/3/04	Canceled	
A6	1/4/04	Canceled	

Table 1. Spirit Trajectory Correction Maneuver Results

All of the maneuvers were executed well within the design requirement for no more than 5% magnitude error (3 sigma). Table 1 shows the results of the four TCMs actually needed and executed for Spirit, and Table 2 shows the results of the three TCMs actually needed for Opportunity. The maneuvers 5, 5X, and 6 on Spirit, and 3, 5, 5X , and 6 on Opportunity were not executed, as the orbit determination results showed that the respective maneuvers were unnecessary. This was good news from a fatigue management perspective. The last few weeks of cruise were very stressful, and the elimination of several critical activities reduced the workload. The ability to target these rovers this accurately was significantly driven by a navigation data type called Delta Differenced One-way Ranging (dDOR). One of the key lessons learned from the cruise operations phase was how powerful this technique actually was.

One key activity not shown in the timeline was associated with each of the maneuvers beginning with TCM 4. This was an EDL parameter update – an opportunity to examine the newly changed trajectory against the EDL timeline (as discussed in section 5.4), and make small software parameter changes to ensure the safety of the vehicle.

TCM	Date (UTC)	Total Delta V (m/s)	Magnitude Error
B1	7/18/03	16.172	-1.40%
B2	9/8/03	0.534	1.10%
B3	11/21/03	Canceled	
B4	1/17/04	0.107	0.60%
B5	1/23/04	Canceled	
B5X	1/24/04	Canceled	
B6	1/25/04	Canceled	

Table 2. Opportunity Trajectory Correction Maneuver Results

4.2.2 Other Cruise Activities

Besides the maneuvers performed to get the rovers to Mars, numerous other activities were performed during cruise. In the following sections both the major activities noted on the timeline and other activities of note are described.

Attitude Control/Navigation Characterization.

During this activity the thruster system to be used for turns and maneuvers was briefly used. This enabled the Navigation team to characterize the basic performance of the system.

Camera checkouts

For the camera checkouts, two dark images were acquired from each camera - this enabled the team to screen out transient signals (such as charged particle trails) from the first image to the second. A second, abbreviated ICO was conducted one month before landing to verify the health of the rover cameras (in particular, the EDL-critical Descent Image Motion Estimation Subsystem camera, DIMES).

Delta Differenced One-way Ranging.

This navigation data type and technique is one of the more powerful in the navigation arsenal. Observations are made of a Quasar, then the spacecraft, then the quasar again by two different Deep Space Network tracking antennas, located at different complexes. These complexes are located at Madrid, Spain; Goldstone, California (USA); and Canberra, Australia. With the information acquired in this way, the spacecraft's angular separation from the known location of the quasar (as viewed from the Earth) is virtually nailed down.

Flight Software (FSW) Load.

Both Spirit and Opportunity were launched on FSW version 7.1. During the cruise period, development and testing continued on version 8.1. This new version contained updates to the EDL capabilities, as well as the final capability to perform surface operations. As such, it was a major effort during the cruise period, not just for the development and test, but also for the three days of file loads and software initialization required to place it on-line for each rover.

Heat Rejection Subsystem Maintenance.

Also not shown on the timeline, but a frequent cruise activity was HRS maintenance. The Heat Rejection Subsystem was designed to transfer heat from the rover's Warm Electronics Box (WEB), deep inside the folded up rover, lander, and the aeroshell, out to the radiators on the cruise stage. This was accomplished via a Freon refrigerant system, using two redundant pumps. The maintenance activity consisted of turning on the backup pump to ensure that the pump received adequate lubrication and was functioning normally.

IMU, RAS, UHF checkouts

During this activity, the Inertial Measurement Unit, the Radar altimeter Subsystem, and the Ultra High Frequency radio subsystem were checked out. The Inertial Measurement Units were located in two places: the inside of the backshell of the entry vehicle to be used during entry and descent, and inside the rover itself to serve as a backup during the EDL events and as the sole inertial attitude measurement during surface operations.

Instrument Checkouts.

Data were acquired from all of the MER science instruments during Cruise Instrument Checkout (ICO) activities. The first checkout, conducted a few weeks after launch, was used to verify post-launch instrument detector and electronics health in the ambient cruise environment. These checkouts consisted of "dark" data acquisition from each instrument, which provides basic information about signal bias levels, noise characteristics, and characteristic detector readout signatures. Performed on each vehicle, these checkouts were intended to serve as the only verification that the instruments had survived the launch event, and that they were ready for Mars operations. The intent was to allow a heads up on any anomalous instrument performance – which was a wise decision. It was during this event that the first indication of a problem with the Moessbauer instrument on both vehicles was detected (see section 4.2.3.4)

Lander Battery Checkout

On board the lander was a set of non-rechargeable batteries, used during the time after the aeroshell has separated from the cruise stage through successful opening of the rover solar panels after landing. A small activity to verify the current and voltage characteristics of these batteries discovered a passivation problem. A subsequent activity was performed just before EDL to de-passivate the batteries (see section

Lander Battery Depassivation.

This activity was performed to eliminate a condition called battery passivation. This is where a coating forms on the cathode of a battery, preventing full voltage and less than normal current draw. As current is drawn from the battery, this layer is broken down and removed. This activity was performed the last day prior to each EDL.

Mini-TES Cross-Calibration of Temperature Sensors.

Not shown on the timeline, but a frequent event during cruise was a cross-calibration for two platinum resistor thermistors in the Miniature Thermal Emission Spectrometer (Mini-TES). These were performed multiple times during cruise, the purpose being to have an accurate cross calibration between two different temperature sensors. Each of the two sensors had been bonded down to the back of the PMA head with two different techniques. These sensors had experienced bonding problems during the pre-launch period Assembly, Test, and Launch Operations period, and it was unknown which technique would be expected to survive the planned thermal cycles during the Martian day and night. These calibration events changed in frequency during the cruise period, but eventually settled down to a frequency of once per month.

Pointing Updates (Attitude Maintenance Turns).

These turns were performed simply to ensure the Earth remained closely aligned with the bore-sight of the medium gain cruise antenna, and the solar panels were aligned to the sun. This is made more difficult in a planning sense, due to the need to consider both the nominal operations plans, as well as different fault cases, including possible fault responses calling sun acquisition.

Solar Array Switching.

Shown on the timeline as various pie shaped segments, another event during cruise was the switching in of additional solar array segments. The solar arrays on the cruise stage were designed to allow us to gradually bring them on line curing the cruise to Mars. The power and thermal balance of the vehicle was modulated through this mechanism, to allow the vehicle to have sufficient power to perform normal (and fault mode activities, if needed), while keeping the temperatures within the allowable operating temperatures. The various solar array stages allowed the power generation to be fine-tuned to 1/8 of the total array capability.

Spin down

One of the first activities after launch was the spin down event. The launch vehicle's third stage spun the cruise stage down to about 12 rpm before separation. After Deep Space Network signal acquisition, the vehicle was commanded to de-spin to the nominal cruise spin rate of 2.0 rpm in increments of 2.0 rpm.

Sun Sensor Side-Head Calibration.

Not shown on the overview timeline, but one of the critical activities pre-EDL, was the calibration of the side heads of the sun sensor. The sun sensor consisted of five heads, two facing axially that were used throughout cruise to provide information on the position of the sun with respect to the spacecraft spin axis, and a set of three heads facing radially out. These radial heads were designed to provide the same sun information, but in the EDL case, where the spacecraft must be oriented with the spin axis towards the entry corridor.

4.2.3 Cruise Problems

Cruise was fairly uneventful, but several concerns did show up. These ranged from a massive solar flare to discovery of thermal design issues.

Each of these took their toll on the flight teams attention and expertise.

4.2.3.1 Record Breaking Solar Flare

The solar flare on October 28, 2003. This was a record breaking solar event, currently assessed as a magnitude X43 with the final intensity still being debated as of this writing. The initial symptoms were a rapid increase in the number of times the spacecraft reported loss of star ID. As charged particles impacted the star scanner, they were each interpreted as a star sighting. These quickly escalated to the point where the spacecraft was unable to distinguish the real stars through the noise created by the flare. At this point, both vehicles dropped out of celestial reference mode and into "sun line" (essentially a lower attitude fidelity mode where the pointing attitude of the spacecraft is established by the sun sensor). The vehicles remained in this state until November 8, 2003, when the vehicles were returned to celestial reference mode.

As a result of this flare, it was discovered that the Mars Odyssey spacecraft had suffered a set of errors in the memory configuration register. This was cleared on Odyssey through the use of a sleep/wake cycle, and the same process was used on each of the twin rovers as well.

Due to continued eruptions from this region of the sun, the proton flux from the flare events did not return to an elevated but reasonable condition until October 8. Additional events continued periodically from this same location, but were not a major problem until November 20. A new set of flare events raised the background level of charged particles at the rover area until November 25. Additional flares of lesser strength continued to occur until the entry, descent, and landing events removed the concern.

4.2.3.2 Thermal Modeling of Thruster Lines.

An additional problem we had was the series of surprises we received in the thermal modeling of the propulsion system lines (ref. 3). The design included a computer control system to monitor and control the heaters on the tank lines. Soon after launch, a flaw in the design manifested itself. The design of the system allowed a sensor on a section of line experience continuous (cooling) hydrazine flow to control not only the heating for that section but also a stagnant section that received no flow. This design flaw was not caught during system level testing, primarily due to the system not containing hydrazine, nor were thrusters fired during the testing.

After that discovery, significant effort by the thermal team was required to analyze the effects of any planned major thruster firings (primarily TCMs), and ensure that thermal limits would not be exceeded.

4.2.3.3 LPSIF SEUs.

Another concern that cropped up was the Lander Pyro Switching Interface (LPSIF), and a series of single event upsets in that part. The LPSIFs provided the end circuits and control electronics for several pyro events that were commanded by flight software during the EDL phase of the mission (e.g., Cruise Stage Separation, Parachute Deploy, etc). The LPSIF is comprised of two completely independent and identical strings of electronics that control the 'A' and 'B' pyro branches (see Figure 1 and Figure 2). Each pyro

I/F string is comprised of two FPGA parts and a pair of FET switches that control the flow of current to pyro relays. All four parts (two per string) are loaded from PROM whenever the PLEM recovers from reset. The parts, once loaded, provided for several pyros that could be fired by software immediately. These were called 'quickfire' pyros. It also provided for a number of pyros that could be set up to fire at some software specified future S/C time. These 'timed' pyros allowed software to set pyros to fire at some default future time and then allowed software to improve the fire time estimate as new data became available. This gave the S/C some resilience to events that might interfere with software's ability to improve pyro event time estimates (for example an unexpected reset).

The parts in the LPSIF are known to be SEU soft. Pre-launch calculations place the expected upset rate at about once every 9 days. This calculation compares very well with the empirical data gathered during the MER cruise phase. The softness of the parts was known and mitigated by the design of the LPSIF. This was accomplished by using hardware to periodically check the configuration state of each FPGA. Whenever a configuration error was detected, the hardware automatically reset the two FPGAs on the side that indicated the error and reloaded the configuration data from PROM. The detection and reload took approximately 10 seconds. During the reconfiguration process, the pyros implemented in the affected string could not be fired. However, the pyro event would still be triggered by the unaffected FPGAs in the redundant pyro string. Failure to fire pyro events from both strings would require an SEU in each string within the same ten second reconfiguration interval coincident with a scheduled pyro event time. Calculations performed pre-launch, and then re-visited pre-EDL, indicated the probability for such an event was extremely remote.

4.2.3.4 Moessbauer Instrument anomalies.

During the early cruise instrument checkout, it was determined that there were two differing anomalies, one for each vehicles Moessbauer spectrometer. During the MER A Instrument Checkout (ICO) on 7/17/03, the Moessbauer Instrument drive velocity error signal during initial cruise checkout was heavily distorted, and reference detector spectrum was correspondingly noisy. The Moessbauer uses a voice coil with a Co-57 source on one end. This voice coil is driven at different frequencies to allow the reflection from the sample of the hyperfine spectral lines to be analyzed. One piece of the information telemetered back to Earth is the difference between the desired drive velocity and the achieved velocity. Subsequent analysis and test showed that there was a subset of drive velocity and frequency settings that would achieve good measurements. It was suspected that there might be an object stuck in the sensor that was impeding the motion of the voice coil.

During the first cruise instrument checkout for Opportunity, the Moessbauer sensor head generated an anomalous reference spectrum: A broad smear, instead of the expected sextet from the reference sample. The error signal, energy spectra, main detector noise levels, parameters, logbook, and temperatures appeared as expected.

No further attempts were made to solve these problems, as it was expected that the landing event might make matters worse. However, in a surprise to all of the engineers and the Moessbauer team, post-landing data from both rovers showed that the Moessbauer anomalies had all corrected themselves. Since that time, instrument data has been remarkably clean, with no anomalies.

5. ENTRY, DESCENT, AND LANDING

The Entry, Descent, and Landing phase of the mission was the most exciting part – both for the flight team members and the world. Described as "6 minutes from hell" by one NASA official, this period was marked by both the most rigorous design and the highest review, and the event had to be entirely autonomous. The one-way light time to Mars prohibited any chance for ground in the loop intervention. The landing site had to be scientifically compelling, and yet within the ability of the landing system.

5.1 Landing Site Selection Approach

The guiding principle behind the selection of the Mars Exploration Rover landing sites was to find the two sites most likely to meet the science objectives of the mission, while being safe enough for landing. The principle could have been to simply land in the two safest places on Mars. However, MER had a very specific science objective, which was to explore environments that may have had conditions suitable for life, i.e. ancient liquid water environments. The determination from orbital data that a landing site both hosted an ancient water environment and has evidence of that environment that the rover's scientific payload could uncover is inherently uncertain and speculative. Hence, two rovers going to two landing sites would significantly reduce the risk of not meeting the science objective. In fact, a major part of the rationale for the addition of a second rover to the project was to reduce the science risk and increase the science diversity by landing at two scientifically distinct sites on Mars. As a result of these considerations, the anticipated science return of the landing sites was paramount, to the extent that two successful landings at sites that both had little or no expectation of evidence of an ancient water environment would be a failure.

The landing site selection was implemented through the following approaches. First, to have an open process involving as much of the Mars scientific community as was willing and able to participate, in order to have not only the broadest acceptance of the landing site selection, but more importantly to have the expertise of that community applied to the scientific and engineering evaluations of the candidate sites. It was certainly in the interest of that community to participate since they would be analyzing the data from these missions for many years to come.

Second, to use all available data on the landing sites from orbital data, including new targeted data from operating orbiters provided as requested and acquired, as well as data from Earth-based observations and modeling in order to understand both the engineering and scientific characteristics of the sites to the greatest extent possible with existing assets.

Third, to develop a detailed simulation of the MER entry, descent, and landing (EDL) system in the relevant environmental models of the

Martian atmosphere and surface through a very close partnership between the EDL engineers and the Mars environment scientists, and to use that simulation to evaluate the landing safety of the candidate sites in order to judge that against the scientific benefit of those sites. Much of the following discussion will focus on the engineering constraints on the landing site selection and the results of the landing site selection process. More details on the scientific aspects of the landing sites can be found elsewhere[5].

5.2 Landing Site Selection Engineering Constraints

The engineering constraints fall into two broad categories, the first concerning the survivability of the EDL event, and the second concerning the quality of the surface mission that would follow a successful landing.

The first category for EDL includes the elevation and elevation variation over three length scales, the rock coverage, predicted wind characteristics at the time of landing, and the ability of the surface to support the RADAR and visible sensing used by the entry, descent, and landing (EDL) system. A key environmental model that itself did not discriminate among the sites, but was required to evaluate the suitability of the site elevation is the density of the atmosphere as a function of altitude.

The second category consists of those factors that determined the quality of a surface mission that would follow a successful landing, with respect to the surface mission lifetime, energy availability for surface activities, rover mobility, and data return. These factors included the latitude of each site with respect to the solar latitude through the surface missions, the nighttime temperature evaluated using the thermal inertia and albedo across each site, the rock coverage, the surface roughness at the scale of the mobility system, and the load bearing capacity of the surface. Additional environmental factors were evaluated that while they did not discriminate between the sites, they were key in determining the suitability of the sites with respect to the required surface mission lifetime and energy availability. These included the dust opacity of the atmosphere, and the rate of dust accumulation on the solar panels. The last factor is the distance between the two selected sites, which could have affected the total data return of the two missions through the orbiting data relays on Mars Odyssey and Mars Global Surveyor. As it turned out, all of the sites were sufficiently spaced, greater than 37° of central angle, that this was not a factor, so long as both rovers were not sent to the same site.

Both categories were evaluated across landing ellipses, which represent the accuracy to which the landing site can be targeted. This was driven by the accuracy of the delivery of the vehicle to the desired atmospheric entry location relative to the landing target, and by the uncertainty in the model profile of atmospheric density. The delivery geometry and uncertainty were the main drivers on the dimensions and orientation of each landing ellipse, and these characteristics varied with the site latitude and landing day (MER-A vs. MER-B). The uncertainty in the density profile did not itself discriminate among the landing sites, but was a significant factor in the determination of the ellipse sizes. Typical landing ellipse sizes were 80 km in length and 12 km in width, oriented roughly East-West in the long direction.

The most basic and limiting landing site characteristics were the allowed latitudes and altitudes. In order to provide the required 90-sol surface science mission, MER-A had to land between 15° South and 5° North, and MER-B had to land between 10° South and 10° North. In order to provide enough time for the EDL events to complete, the altitude had to be below -1.3 km MOLA reference. This immediately defined portions of a band in which landing sites could be identified. The -1.3 km was science driven, since that was the altitude of the highly desirable Meridiani Planum site. This in turn drove the EDL design to accommodate that altitude.

The rock coverage was characterized using thermal inertia data from the Viking orbiter IRTM instrument, which had a resolution of about 1°. This provided two to four samples of the rock coverage for each ellipse, with a sample expressed as a percentage of the surface area covered by rocky material. The initial constraint was 20% rock coverage, derived from Mars Pathfinder. The rock coverage of the specific sites was used later as an input to the simulation. The rock coverage percentage was used to generate a rock size distribution for that level of coverage, based on Earth analogues and validated in one instance through observations by the Mars Pathfinder lander. The probability of hitting large rocks during the landing event is a key factor in determining the success of the landing. Also the rock coverage determines the trafficability of the terrain for roving, but that was not a discriminator among sites, since if they were safe enough for EDL, they were trafficable with respect to rock coverage.

Thermal inertia measurements from Viking and Mars Global Surveyor (MGS) was also used to model the nighttime temperatures at the landing sites for the surface missions, and to qualify the sites for adequate RADAR signal return and load-bearing. In addition to the thermal inertia, the albedo of the surface, as measured by the Odyssey THEMIS instrument, was a required input for the nighttime temperature models. The diurnal temperature cycle towards the end of the required 90-sol surface science mission, especially the profile of the nighttime temperatures, was a key determinant of surface lifetime of the rover.

One of the most significant data sets used was the MGS MOLA instrument, a laser altimeter. This provided detailed topography for all of the sites, and was used to establish the site elevation, site slopes at the 300 meter scale and greater, and to provide detailed topography for mesoscale wind modeling. The slopes at around the one kilometer scale were required to be low enough that the airbag system would not continue to roll and especially not accelerate while rolling downhill.

Mesoscale wind models are very compute-intensive 3-D simulations of the Martian atmosphere over small regions of the planet. They took as input the global circulation models for boundary conditions, assumed atmospheric opacity, atmospheric temperature profiles measured by the MGS TES instrument, and the underlying topography from MOLA. From these models a great deal of information can be extracted concerning the motion of the atmosphere as a function of the local solar time, the altitude, and the location. Though these models necessarily had a high degree of uncertainty due to a lack of significant direct wind measurements on Mars, they provided a means to assess and compare the landing sites with respect to wind. Horizontal winds and wind gusts represented a threat to EDL due to induced horizontal velocities on impact.

In addition to the MOLA topography, an even finer measurement of surface slopes and roughness was provided by MGS MOC stereo images over small subsets of the landing ellipses. Slopes with length scales as short at 10 meters could be resolved, which was required for the evaluation of the bounce characteristics of the airbag system on the terrain, as well as the performance of the EDL RADAR altimeter with respect to changes in the measured altitude and descent rate due to terrain variations in the last few hundred meters of descent. These stereo data sets were collected over representative terrain types, which were themselves classified using MGS MOC and Odyssey THEMIS visible data that covered the candidate landing sites. The small-scale slopes and their variation was used to model the bouncing of the airbag system, which depending on those slopes could transfer energy between the horizontal and vertical velocity components. Since the vulnerability of the system varied across those components, the small-scale slope modeling was a significant safety discriminator between landing sites.

Earth RADAR observations of the landing sites using the Deep Space Network provided corroborating information on the RADAR return at the sites, but more importantly they provided a much finer scale roughness measurement of features at decimeter scales that could not be obtained by other means. This data type was important to assess trafficability of the rovers at the scale of the rover wheels, which are 26 cm in diameter.

5.3 Landing Site Selection Timeline

The MER project began in May, 2000, and the landing site selection process commenced almost immediately with preliminary landing site engineering constraints completed in August, and first set of 155 candidate landing sites in September. These candidates were identified using only the engineering constraints to provide the broadest possible set for science consideration. These candidates were distributed to the science community, and the first landing site selection workshop brought that community together in January 2001. At that workshop, seven high priority and 17 medium priority landing sites were selected for more detailed evaluation and observation based on their scientific potential. The Mars Global Surveyor and Mars Odyssey orbiters began targeted observations of the high priority sites. More detailed engineering constraints were developed as the design of the MER flight system and navigation capabilities progressed. These updated constraints were provided in July 2001, and a second landing site workshop was held in October. Due to the investment required in new orbital observations as well as analysis by the project, the workshop was requested to select four landing sites for continued consideration, with the final selection of two sites expected, at that time, to occur one year later in September of 2002. The October 2001 workshop resulted in the selection of Meridiani Planum, Melas Chasma, Gusev Crater, and

Athabasca Valles as the four primary sites, with Isidis Planitia and Eos Chasma identified as backup sites in the event that later engineering evaluations ruled out some of the primary sites. One promising site, Gale Crater, dropped out of contention at that time due to the fact that the landing ellipse would not fit in the crater. Significant effort to reduce the size of the ellipse did however permit the larger Gusev Crater to remain in the running.

By early 2002, the first results from the mesoscale wind models were becoming available. Unfortunately, they showed very high winds in the early afternoon through the Valles Marineris canyon, over the Melas Chasma and Eos Chasma sites. These winds could easily peak higher than the MER EDL system could handle or compensate for, and so those sites were eliminated after the April 2002 landing site workshop. This left three prime sites, Meridiani Planum, Gusev Crater, and Isidis Planitia, which had been promoted, and Athabasca Valles, which had been demoted to a backup site. Preliminary analysis of Earth RADAR observations showed anomalous roughness at decimeter scales in parts of Athabasca. This essentially put Athabasca on probation until the analysis could be completed and conclusions could be drawn about the anomalous RADAR return. In June 2002, it was concluded that untrafficable terrain was a plausible interpretation of the data, and so the Athabasca site was taken out of the running.

At that time, only three landing site candidates remained, and only one, Meridiani Planum, appeared to have a high certainty of being certified as safe. Both Gusev Crater and Isidis Planitia had higher winds, and it was possible that final airbag testing would rule out those sites as well. As a result, the project requested a fourth landing site candidate be identified that was "wind-safe", that is, would be likely to have low enough winds after more detailed mesoscale wind modeling to be certified as safe. That site was to be selected without concern for the science return of the site, as a backup for the second rover. In August of 2002, the Elysium Planitia site was identified, which indeed looked like it would meet the safety criteria, but unfortunately it appeared to consist largely or entirely of volcanic deposits unlikely to bear any evidence of an ancient water environment. This was the only time in the process that a site was considered that did not have the potential to meet the science objectives. Also the latitude constraint was relaxed slightly to permit this site at 11° North.

In January 2003 the fourth and final landing site selection workshop was held. During this meeting the landing sites were prioritized for science value and testable hypotheses were developed for each site. This was followed by a project science team meeting, in which Meridiani Planum and Gusev Crater came out as the highest priority sites, with Meridiani Planum holding the majority as the site to send a single rover if events required that decision, such as a launch failure of one rover. Isidis Planitia was third, and Elysium Planitia was a very distant fourth, due to its geology. The Meridiani Planum and Gusev Crater sites were highly complementary in the lines of evidence suggesting a liquid water past. Meridiani's evidence was mineralogical, based on the detection of the mineral hematite by the MGS TES instrument. On the other hand, Gusev's evidence was morphological, indicated by a very large channel that flowed into what appears to be a crater lake. This provided a very good level of scientific redundancy.

As the airbag testing and simulations progressed through early 2003, it appeared that while the Meridiani Planum site was safe enough, Gusev Crater was borderline. As a result, in April 2003, the project recommended to NASA Headquarters that MER-B be sent to Meridiani Planum, and that MER-A be retargeted to either Gusev Crater or Elysium Planitia depending on further airbag testing to be conducted while the rovers were enroute to Mars. This recommendation was accepted.

Both rovers were launched successfully, and the first maneuvers targeted MER-A halfway between Gusev and Elysium, and MER-B directly to Meridiani. Due to favorable launch vehicle accuracy, MER-A had sufficient propellent to target directly to Gusev Crater at its second maneuver in July 2003, and still have enough propellant to retarget from there to Elysium in early November. While MER-A was on target to Gusev Crater, the airbag testing and analysis was completed by early September 2003. Significant airbag capability was revealed in those tests, and Gusev Crater was certified as safe and NASA HQ approved the final landing site selection.

Though the landing sites were selected, some of the landing site selection analyses continued. In December 2003, less than a month before the MER-A landing, a significant regional dust storm developed over the Ares Valles region of Mars, which affected the dust opacity and temperatures of the atmosphere over the entire planet. This required a re-evaluation of the atmosphere density models and a re-analysis of the atmospheric entry targeting. While the target on the surface remained the same, the new atmosphere model changed the atmospheric entry point. These changes were incorporated into the final maneuvers for MER-A.

On January 3rd, 2004, MER-A landed successfully in Gusev Crater. The expected environments were seen, and the EDL systems designed to compensate for wind were activated and successfully countered the winds at that site. On January 24th, 2004, MER-B successfully landed in Meridiani Planum, and again the environments were as predicted, including the very low winds with the EDL wind compensation not activated. The landing site selection process was successful in predicting the characteristics of the two landing sites, and to date one of those sites, Meridiani Planum, has provided ample evidence of an ancient liquid water environment on Mars. Such evidence has not been found at Gusev Crater, justifying the desire for science redundancy through two landing sites. However the MER-A mission is still ongoing, and so may yet provide more evidence of an ancient water environment in a different locale on Mars.

5.4 EDL Overview & Timeline

At first glance at the system level, the MER cruise and EDL system appears nearly identical to the MPF EDL system. Key among the similarities is:

1. Passive 2 rpm spin-stabilized cruise and entry.
2. Active thermal control using circulated Freon, vented prior to cruise stage separation.

3. Nearly continuous X-band "communication" to the Deep Space Network (DSN)
4. Pre-entry cruise stage separation
5. 2.65 m, 70 degree cone Viking-like entry aeroshell using SLA-561V ablative material
6. Single-string 20 MIPS RAD2000 (in rover) controlling both cruise and EDL
7. Acceleration-based software-controlled parachute inflation timing
8. Viking & MPF derived disk-band-gap parachute
9. Ballasted heatshield separation
10. Lander-backshell separation and deployment on 20 m bridle
11. Single RADAR altimeter for software assessment of altitude and descent velocity
12. Three canted solid rocket motors mounted inside the back of the entry aeroshell
13. Software-controlled airbag inflation, Rocket Assisted Decelerator (RAD) firing and timed bridle cut.
14. Multiple airbag bounces
15. Computer-controlled airbag retraction and vehicle righting during lander petal opening.

Despite the many hardware similarities there are significant differences:

1. The single computer controlling EDL is inside the rover
2. A large (40-60 deg) autonomous turn to the entry attitude is required
3. 41% higher entry mass than MPF (832 kg for MER)
4. 28% lower entry velocity than MPF (5.7 km/s for MER)
5. Steeper entry flight path angle (-11.5 deg for MER vs. –14.2 for MPF)
6. Daytime landing resulting in lower atmosphere density and higher winds than MPF
7. A 1.3 km higher landing site elevation than MPF (-1.3 km MOLA for MER at Meridiani)
8. A 6-DOF Inertial measurement Unit (IMU)-based propagation of both the rover and the backshell attitudes.
9. 28% larger DGB parachute (8.9 m diameter for MER)
10. Descent camera as part of the Descent Motion Estimation System (DIMES)
11. Image processing software that used the descent camera's images to calculate proper horizontal velocity (DIMES)
12. Three small (1200 N-s each) transverse impulse rockets (TIRS) in the backshell for horizontal velocity mitigation
13. 87 % larger RAD motors than MPF (33000 N-s per canted rocket for MER)
14. Attitude control software that used backshell IMU data at the time of RAD ignition and DIMES inputs to fire TIRS and thus control touchdown horizontal velocity
15. Strengthened MPF 5m diameter Airbags (6 abrasion and 2 gas containment layers on MER)
16. Descent UHF radio relay link to Mars Global Surveyor

Autonomous on-board software controls most mission-critical events on the vehicle from the time that the EDL software is initiated 5 days before landing until hours after it has landed once the vehicle is in a power-safe state on the surface of Mars.
As the one-way light time for signals to reach earth at arrival was nearly ten minutes, there is no way that computers or people on Earth can intercede in any way during EDL. Therefore all mission-critical actions are handled by software executing on the single computer inside the rover.

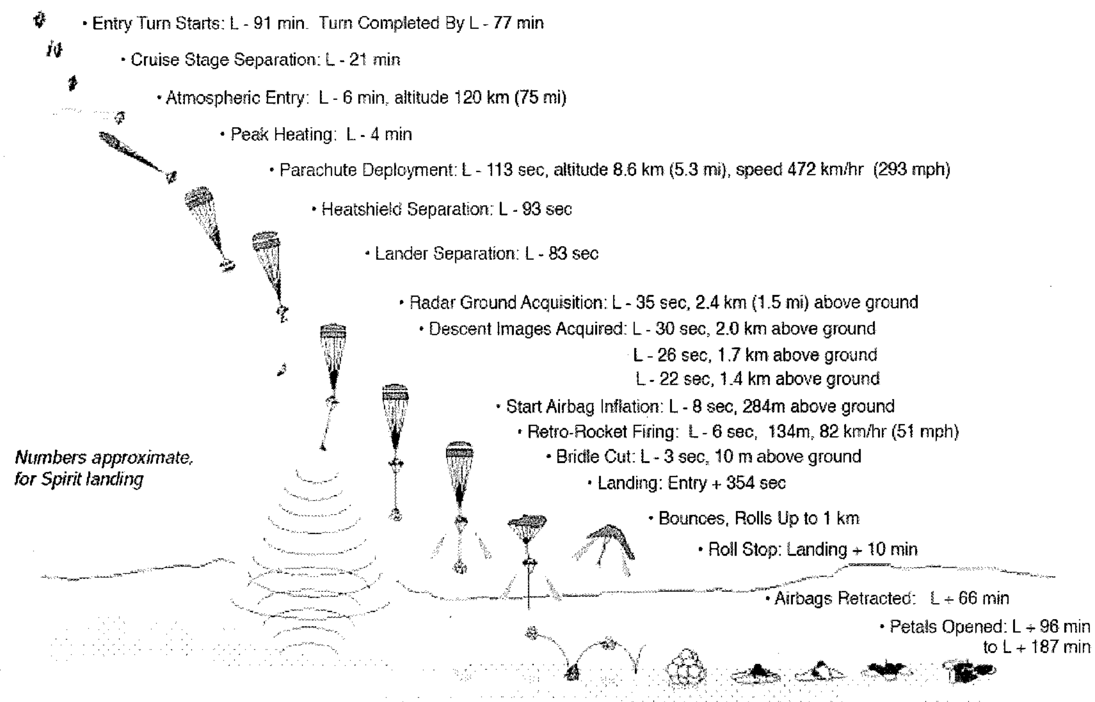

Figure 3. Entry, Descent, and Landing Timeline

While the EDL software is enabled days prior to landing, the first autonomous event occurs about 2 hours before entry when the software instructs the cruise attitude control system to begin preparations for the turn to entry (see Figure 3). Unlike Mars Pathfinder, the posigrade approach geometry necessitated about a 22 deg turn for Spirit and a 42 deg turn for Opportunity. The turn begins about 1 hr and 25 m prior to entry and completes in about 20 minutes. At 40 minutes prior to entry, the Freon coolant that had been circulating about the cruise stage and into the rover since launch is vented into space via a vent tube along the spin axis of the vehicle. At 15 min prior to entry, the cruise stage, having completed its mission to provide power, thermal, and attitude determination and control for the vehicle is separated via three pyrotechnic separation nuts. The 2-rpm spin rate of the vehicle used in cruise is maintained throughout entry for stability. Unlike the Viking missions that used attitude control thrusters, both MER and MPF used spin stabilization for attitude control during hypersonic entry.

The entry point is reached when the vehicle is 3522.2 km from the center of Mars or about 125 km above the surface where the vehicle is traveling at 5.7 km/s. Two minutes later the vehicle experiences about 5.8 Earth Gravity's of peak deceleration. In the following two minutes the vehicle decelerates to 0.4 km/s. The on-board software uses IMU accelerometers to determine the optimum time to mortar-launch the parachute so that the dynamic pressure at the time of parachute inflation is under a ground-specified value. In the case of Spirit, this value was 725 Pa and in the case of Opportunity, this value was 750 Pa. The vehicle is about 6-7.5 km above the surface when the parachute is inflated.

Twenty seconds after parachute deployment, the ablative heat shield is separated from the backshell via firing of two sets of 3 pyro separation nuts. Ten seconds later another two sets of 3 pyro separation nuts separate the lander tetrahedron containing the rover from the backshell allowing the Zylon bridle to be deployed to its full 20 m length. In this configuration, the single RADAR altimeter is mounted on the underside of the lander. The EDL software uses it to estimate the terminal descent rate and altitude. The RADAR locks up on the ground at 2.5 km altitude while descending at approximately 70-80 m/s. The descent rate and position is used to estimate when to initiate the firing of the three main Rocket Assisted Decelerators (RAD) rockets as well as to time the cutting of the bridle some 8 m above the ground.

Also in this configuration, the descent UHF antenna on top of the lander tetrahedron is in the deployed position allowing high rate data 8 kHz to be transmitted to the Mars Global Surveyor spacecraft that had been strategically timed to be overhead during the decent phase. In addition to this data the software is sending very low rate (0.1 Hz) multi-phase-shift-keyed (MFSK) modulation of the vehicle state to the Deep Space Network in Canberra, Australia and Goldstone, California via the rover's X-Band low gain antenna stack.

As the vehicle descends under the parachute, three images are acquired from a small camera mounted near the RADAR antenna when about 2000, 1700 and 1400 m above the ground. Lander orientation estimates are used with these images using attitude-corrected 2-D image correlation techniques to determine ground-relative velocity at these altitudes. This system is called the Descent Image Motion Estimation Subsystem (DIMES).

The DIMES horizontal velocity is then propagated via the IMUs along with estimates of the relative orientations and positions of the backshell and lander. Just prior to RAD firing the software uses these estimates to determine what best combination of the three small horizontally-mounted Transfers Impulse Rockets Subsystem (TIRS) rockets should be fired. TIRS software attempts to minimize the horizontal velocity the system would have attained at the time of bridle cut due to the combination of RAD swing angles induced by wind shear and actual horizontal wind gusts determined by DIMES.

The two sets of two airbag tie downs are cut at 2.25 and 2.125 s prior to RAD ignition. The three airbag gas generators are activated 2 s before RAD ignition. The TIRS rockets are fired 0.2 or 1.1 s into the RAD burn (depending on the required magnitude of the velocity correction). The RAD rockets fire at about 120 m for 2.5 to 4 s before the bridle is cut approximately 8 – 10 m above the ground. Once cut, the RAD rockets continue to fire for up to 2 s, launching the backshell and parachute up and away from the bouncing lander.

The landing load capability of the vehicle and the airbags is as high as 40 gees although the most probable load was expected to be about 10-15 gees. The vehicle then bounces as much as 1 km before rolling to stop.

After at least 10 minutes from landing, the vehicle begins to retract the airbags. Near the point that the uppermost bags are retracted, six separation nuts that restrain the petals are fired and one of the four rover tie down cables are released. Next, the side of the lander tetrahedron that is on the ground is opened via high torque petal actuators that right the vehicle. Part way into petal opening, the remaining petals are opened leaving the lander fully open and the rover exposed and upright. The final critical action is for the rover's solar arrays to be pyrotechnically released and opened.

The rover is now safe to survive its first of many evenings on the surface of Mars. All told, 55 pyrotechnic devices (of which 37 are fired before hitting the ground) are fired in the few hours starting with the Freon venting until the solar arrays are open.

5.5 EDL Operations Testing

From July through the first week of December 2003, the MER project was in a continual state of getting the software and the people ready. In particular, the EDL and cruise team members found themselves learning how to prepare the processes. There were four full-up project-wide operational readiness tests (PORTs) that included EDL operations.

The first was PORT 2 in July 2003 whose purpose was to:

1. Demonstrate MER/ODY/MGS UHF strategic planning process
2. Validate the very tight TCM-6 preparation timeline
3. Validate post-TCM-6 EDL Parameter Update Timeline / Process
4. Validate CMSA EDL Joint Operations with Active Cruise Mission
5. Exercise the EDL Doppler and Range Predict Delivery Process to the DSN
6. Practice MGS data flow to MER

While these objectives were met, it was found that there were several areas where there were fundamental process design problems. In particular we found that the many constraints on the maneuver design process resulted in the later trajectory correction maneuvers decompositions that resulted in huge ground track swings that brought the instantaneous impact points far from the target area. Further relaxation of the constraints (entry flight path angle and time of arrival) resolved this problem.

The team also discovered the "fog of EDL" where it is very difficult to discern what is happening especially at very low data rates. With further practice the team learned how to extract and monitor only those measurements that were critical for assessing the vehicle.

The second test was PORT 4/5 (Sept 2003) whose purpose was to:

1. Practice an improved TCM-5 & 6 de-composition process
2. Practice the nominal TCMs 5 & 6 and get the teams working together.
3. Practice getting the complex DSN configuration setup and working via the testbed.

These objectives were also met however the team learned that we had not properly thought through the purposes for TCM-6. It was being overloaded to support both a spacecraft anomaly during TCM-5 at E-2 days and it was also being used to be the last possible maneuver. Further debate suggested that a TCM-5X at E-1 day would suffice for a spacecraft TCM-5 backup window and that TCM-6 could and should be pushed as late as possible so that the effects of the gravity well could be detected in the navigation data.

At this PORT the team elected to run the test on the CETB (including a spare Small X-Band Deep Space Transponder) and pipe the data via an optical fiber to the DSN's "development station" called DTF-21 near JPL. This trick allowed members of the EDL communication team to be able to detect the 0.1 Hz MFSK modulation using the intended tools. However the 9 minute light time was not be present, so it was decided to also run EDL on the surface system testbed (SSTB) 9 minutes later which allowed the flight team in the cruise mission support area (CMSA) to see telemetry as well as UHF data in the flight timeline. This trick backfired on thee team however when due to a setup error, the CETB suffered a short reset early in the timeline that was seen in the DSN data, but not seen in the digital data from the SSTB (despite having been initialized nearly identically). This confused the team (including those in charge) when it was reported that a reset had occurred in one data set but not the other.

The third test was PORT 6 in Oct 2003. Its purpose was to fix the errors that had occurred in prior PORTs and also provide an opportunity for the team to correct a serious anomaly in the approach phase. The anomaly selected was insidious. In this failure scenario during the normal execution of a lateral burn portion of TCM-5 a spin thruster became stuck on. It had become stuck after a micrometeor had impacted the +X thruster cluster. The meteor had also resulted in a very tiny leak in the Freon coolant system that created a nearly continuous small force after the TCM. Once the thruster became stuck on, it was up to the team to make a quick decision on whether to let the on-board fault protection catch the error and stop the maneuver or to attempt to stop the maneuver via emergency contingency commands. The team chose to stop the maneuver leaving the spacecraft in a rather poor power and communication attitude. The team had 24 hrs to recover including coming up with a new maneuver design that took into account the effect of the spin thruster over burn. It took almost 6 hours for the team to diagnose the thruster failure and, later, the Freon leak. It took about 12 hrs for the navigation team to notice the trending of the small forces and correlate that data to the observed Freon pressure decay rate. In short the team recovered and landed the vehicle safely despite many challenges that occurred during this simulated mission.

The final test was PORT 10 whose purpose was to correct prior errors and to follow through with a detailed MER-A EDL reconstruction process, including EDL Red Team review of the results and processes.

In July 2003, it was pointed out that since there was only 3 weeks (nearly to the hour) between MER-A and MER-B landing, if there was a catastrophic failure of MER-A there will be little time for NASA to convene a failure review board prior to the second landing. Instead the team suggested that the MER-A failure review board be convened prior to landing. This team was formed from JPL, NASA and university members and became known as the EDL Red Team. This team was chartered also with the task of approving the EDL reconstruction processes and certifying that the MER team was ready to conduct a quick a thorough investigation into the details of MER operation and to make detailed recovery recommendations within 14 days of landing. PORT-10 gave the EDL team the chance to implement the process and develop the tools. It also gave the EDL Red Team an opportunity to critique the process.

By the second week of December 2003, since launch ten EDL-related reviews had been held, four full-length one-to-two week operational readiness tests had occurred and many hundreds of hours of testing had taken place. The team was as ready as it would ever get.

5.6 EDL Execution

After all that testing described above, the team was ready. On January 4, 2004, MER-A (Spirit) made history. At 4:58 pm PST, only three and half hours before entry, the commands to enable the EDL pyro timers were executed. At 6:15 pm in the Cruise Mission Support Area after the pre-EDL team readiness poll, the mission manager played a familiar "wake up" song for the EDL event entitled "Don't worry. Be Happy" by Bobby McFerrin. At 08:25:36.5 pm PST, watched by literally tens of millions of people, the lander made first contact with the surface of Mars. Nine minutes and 28.5 seconds later the signals arrived at earth at 08:35:04.5 pm PST. Confirmation that the rover had survived landing did not come for another 18 minutes later when the rover began X-Band transmission to the 70 m station at Canberra Australia using a small patch antenna on the bottom of the lander that had a better view of earth than the rover low gain antenna stack. The work of getting the vehicles off the lander was still in front of everyone, but a few hours were spent celebrating the occasion.

Three weeks later, on January 21, 2004, MER-B (Opportunity) repeated the feat. Mankind had successfully placed two more vehicles from Earth onto the surface of another planet.

6. Extended Mission Planning

As has been previously mentioned, the surface operations portion of this mission is described elsewhere. Needless to say, after its successful conclusion, NASA allowed an immediate extension to the mission through September 2004. During this short initial extension, plans have been laid for another, more ambitious extension. This mission will be designed to further the work of the earlier explorations of the two rovers, but will allow for the lessons learned during the preceding nine months, as well as increased confidence in the capabilities of the two vehicles and the flight team that operates them.

While still in negotiation, this new mission will continue to keep the rovers working to return the highest science return for the taxpayer's money. No guarantee can be offered concerning their future lifetimes, as they have already exceeded their expected lifetimes by several factors, but the engineering design used for these rovers is solid and conservative. They have already given Humanity their money's worth, but more will come.

7. Summary

As can be seen from the preceding, the successful landing on Mars of the twin rovers required an immense amount of effort and teamwork, simultaneously operating the mission and preparing for the next phase (including software development for surface during the cruise phase). Accomplished in an unbelievably short three and a half years from conception to landing, this project stressed the abilities and resources of NASA and the Jet Propulsion Laboratory. Although it looked smooth and easy on television on landing day, the success of the mission serves as a testament to the dedication of the teams involved in designing, building, testing, and operating the twin vehicles. The amazing science results coming out of the mission will be studied for decades.

8. Acknowledgements

The authors would like to acknowledge the assistance of the following personnel, without whose aid this paper would not have been possible: Erik Bailey, Robin Bruno, Richard Cook, Joy Crisp, Louis D'Amario, Jim Donaldson, Polly Estabrook, Kirk Fleming, Rob Grover, David Kass, Phil Knocke, Paul Laufer, Wayne Lee, Jan Ludwinski, Justin Maki, Miguel San Martin, Robert Mitcheltree, Glenn Reeves, Ralph Roncoli, Adam Steltzner, and Jason Willis,.

The authors would also like to acknowledge the efforts and dedication of the Mars Exploration Rover Project personnel – the Team. Without their efforts, the missions as described would not have happened. Thank you!

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics Administration.

9. References

1. J. Erickson, et al, "Mars Exploration Rover Surface Operations", IAC-02-Q.3.1.03, October 2002.
2. J. Trosper, "Mars Exploration Rover Surface Operations", IAC-04-Q.3.a.05, To be presented, October 2004.
3. K. Novak, et al, "Thermal design and Flight Experience of the Mars Exploration Rover Spacecraft Computer-Controlled, Propulsion Line Heaters, SAE 2004-01-2412
4. L. D'Amario, "Mars Exploration Rovers Navigation Results", AIAA-2004-4980
5. M. Golombek, et al, "Selection of the Mars Exploration Rover landing sites", JGR, Vol 108, No. E12, 8072