THE MARS EXPLORATION ROVERS: HITTING THE ROAD ON MARS

Nagin Cox

Jet Propulsion Laboratory
National Air & Space Administration/California Institute of Technology
Email: nagin@jpl.nasa.gov

Abstract: Since the beginning of time, people have been fascinated by Mars. From the earliest mission to now-Mars has been (and is) a challenging destination. The Rovers were developed at a breakneck pace in 3 years and landed successfully on Mars in January 2004. This paper will discuss how the Mars Rover mission fits into the overall Mars Program and NASA's program of planetary exploration. Building the rovers in such a short time period created some difficult design challenges that were mainly schedule driven. In addition, it will cover the process of selecting the rover landing sites as well as the engineering challenges faced in the entry, descent and landing process. The rovers have a great deal of autonomous control ability on the surface and the process of developing and testing those was part of the challenge of doing this in 3 years.

Keywords: cameras, interplanetary spacecraft, obstacle avoiding, operations, system engineering, autonomous vehicles, training

1. INTRODUCTION

Every 26 months, the distance between Mars and the Earth is at a minimum. The NASA Mars Robotic Program takes advantage of this favorable geometry to send an interplanetary robotic explorer each time one of these "closest approaches" occurs.

2003 was an unusually good time to visit Mars since Mars and the Earth would be closer together than they had been in thousands of years. Nations from all over the world were sending robotic ships to Mars and NASA planned to be a part of that international fleet. The question was how. One approach was to stay away from the difficult task of landing and put an Orbiter around Mars. Instead, NASA chose to try another landed mission by returning to the proven methods used in the Mars Pathfinder mission in 1997.

[*)] This paper was published in the 16th IFAC Congress preprints under the written permission from NASA.

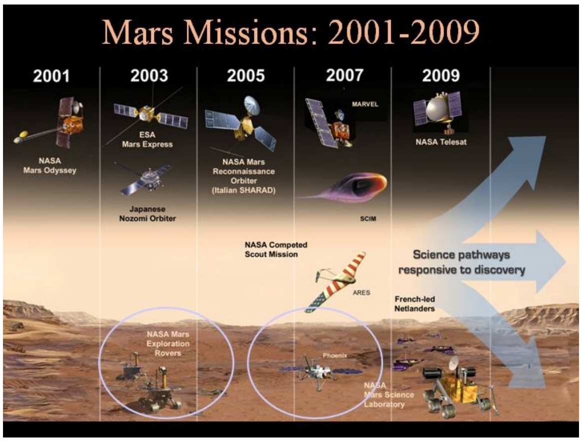

Fig. 1. Possible Mars Missions every 26 months

The Mars Exploration Rover (MER) mission was born as a resurrection of the best of Mars Pathfinder landing system with an updated Rover capable of covering much greater distances than Sojourner (the Rover on the MPF mission).

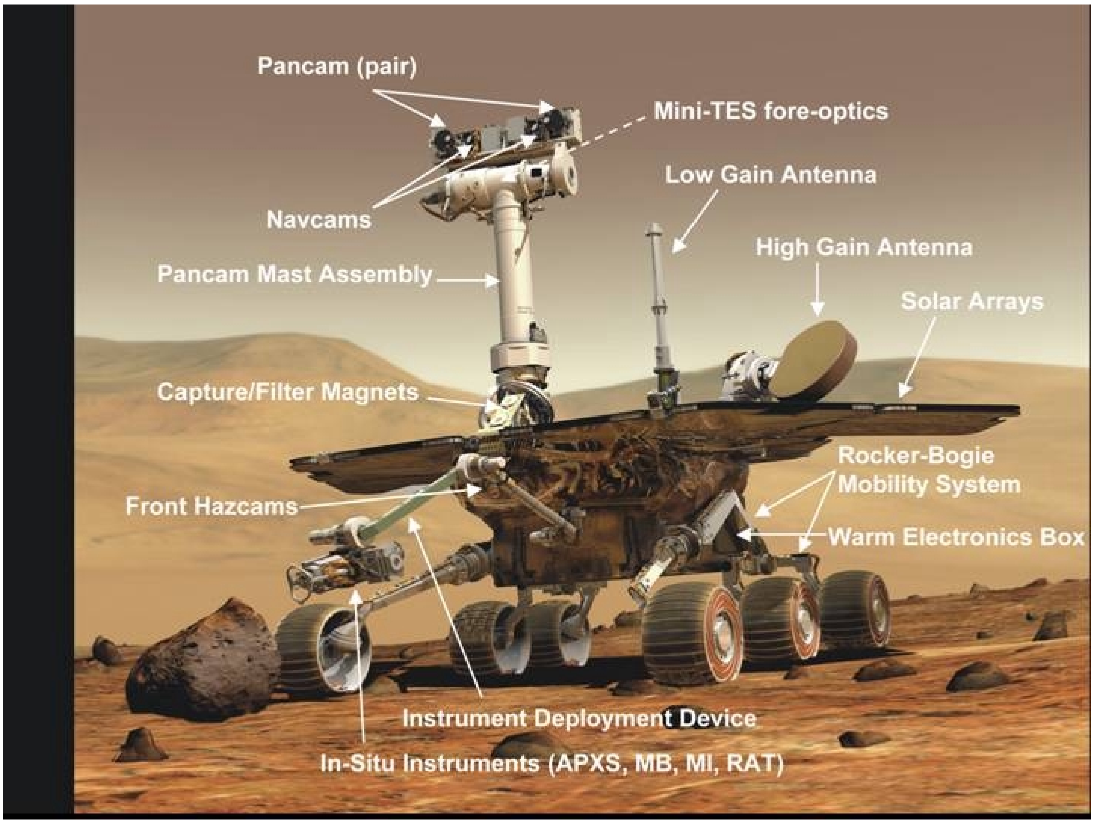

Fig. 2. The Rovers as Robotic Geologists

The Mars Exploration Rover (MER) mission was born as a resurrection of the best of Mars Pathfinder landing system with an updated Rover capable of covering much greater distances than Sojourner (the Rover on the MPF mission).

2. MISSION OVERVIEW

MER had two types of objectives- exploration and science. The science goal was to find evidence of past liquid water on Mars. The exploration goal was to increase robotic mobility on other planets by developing a rover that could traverse multiple meters in one day.

In Mars exploration (and across the rest of the solar system for that matter), the driving theme behind the science goals is “follow the water”. Water is required to sustain life as it is currently known. Thus, the science objectives follow a clear theme of orbiting missions looking remotely for evidence of water based processes and then landed missions that focus on specific areas that orbiting assets have identified.

In the end, the major requirements produced the equivalent of a robotic geologist. A human geologist would locate a rock of interest, move over to it, and examine it with tools at the end of the human arm. Similarly, the Mars Rovers get large scale views of their surroundings and send those back to the earth

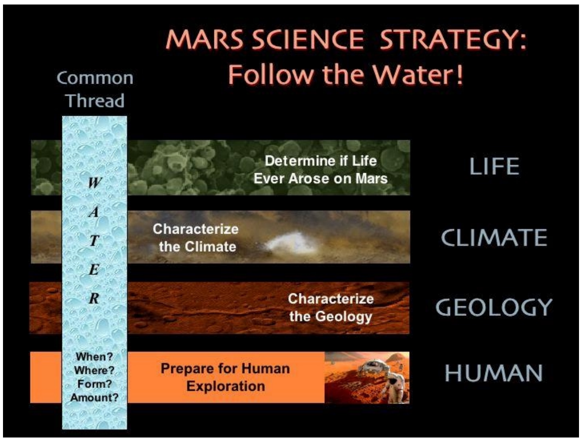

Fig. 3. Follow the Water

where scientists would choose interesting targets. Then those targets would be sent to the Rovers who would maneuver over to the rocks and examine them with a suite of tools at the end of a robotic appendage called the Instrument Deployment Device (IDD).

The Rovers use solar arrays for power as well as a battery system that assists in keeping the Rovers warm at night. The critical vehicle components are housed in the Warm Electronics Box which is designed to stay warm during the cold Martian night. The Rovers can use multiple methods to contact the earth including a low-gain or high-gain antennae as well as a UHF antennae that allows the rovers to send and receive data through an orbiting asset around Mars.

The payload suite consists of multiple cameras used for navigation as well as scientific data gathering. In addition, the rovers carry a 5 degree of freedom robotic arm with a turret at the end containing two spectrometers and a camera as well as a Rock Abrasion Tool (the RAT) used to shave off the top layer of a rock and allow examination of what lies beneath. Much of this suite of instruments was already partially developed for use on the cancelled Mars 2001 lander.

The Navigation Cameras are a stereo pair with a large field of view for wide angle images used for navigation. The Panoramic Cameras with color filter wheels are mainly used for science and have a smaller field of view than the navigation cameras

3. DEVELOPMENT

One of the benefits of using the Mars Pathfinder approach was to be the generous use of "heritage" – i.e. reuse of existing designs or hardware. In the end, very little heritage remained. The Rover was just too new and too innovative in its goals and it was a large challenge to find a way to finish this task in 3 years.

The schedule was so tight that the mission spent only 6 months in preliminary design versus the years that many missions spend in that phase. The requirements of the Rovers started to stretch power and mass constraints significantly

Since the launch date of the rovers could not change, the system engineering teams had to change. There was an initial period where MER was organized like more traditional engineering projects, however, it soon became clear that the task was too large for this kind of hierarchical organization. Instead, the team had to split up into three parallel development teams focused on different phases of the mission.

One team was responsible for building and testing the hardware, software, and ground systems responsible for transporting both rovers to Mars. Another team focused solely on the Entry, Descent, and Landing (EDL) phase- the 6 minutes required to get from the top of the Martian landing to first contact with the surface. The third team focused on the rovers themselves- getting them off the lander and into operation on the surface of Mars. This parallel approach was absolutely necessary to focus the team on one task and get it done.

The design work also had to be optimized to be done at the most efficient time –i.e. when the design of the vehicles was mature enough for work to be able to be done once and not have to be re-done due to design changes. These were the first robotic vehicles to require this level of on-board autonomy. Thus, how the rover was going to be operated was critically tied to how the on-board autonomy worked.

These rovers would be non-deterministic.-a new concept in operations. The ground-based operators would not know where the rovers would be and what their surroundings would be until the end of a Martian day. The earth would not be in contact with the rovers while they were exploring so the rovers had to take care of themselves. Thus, traditional software protection algorithms would not be sufficient. The rovers had to be able to deal with unusual circumstances and continue to operate. Time was of the essence on a mission with a limited lifetime mission like this.

This resulted in The "Fault Protection" software for MER, or the algorithms that usually stabilize the spacecraft in case of a problem until the ground crew back on earth can intervene, being developed after the core autonomous capabilities for the vehicles had been designed.

The other approach, which worked well on the innovative Mars Pathfinder mission was to get into the test phase of the mission as soon as possible. The MER mission had the advantage of being "hardware rich" –meaning multiple platforms for testing. This allowed us to begin testing as early as possible. This occurred at the breadboard and functional testing level as well as being the driving principle behind the Assembly & Test phase as well. The mission had two rovers that were going to Mars and both those Rovers could be used as independent test platforms to complete the necessary suite of testing. Having the two rovers to test with helped the schedule somewhat but in the end it was still necessary to work double shifts for an extended period to meet the launch date.

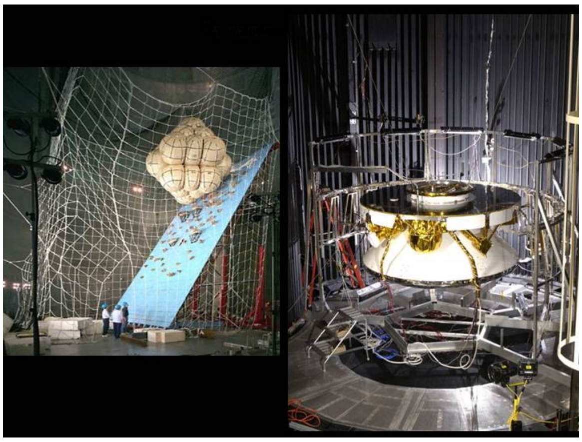

Fig. 4. Test early on as many platforms as possible

In the end, all aspects of the development came down to a few clear guidelines: get the best people, focus them on the 80% of the work that matters and start testing early to find the things that might have been missed.

4. EARTH TO MARS OPERATIONS

The first rover was named Spirit and launched on June 10, 2003 and the second rover named Opportunity launched on July 7, 2003.

The rovers were now flying in the configuration needed to get to Mars- i.e. the rover was tucked away inside the aeroshell that would deliver it to Mars.

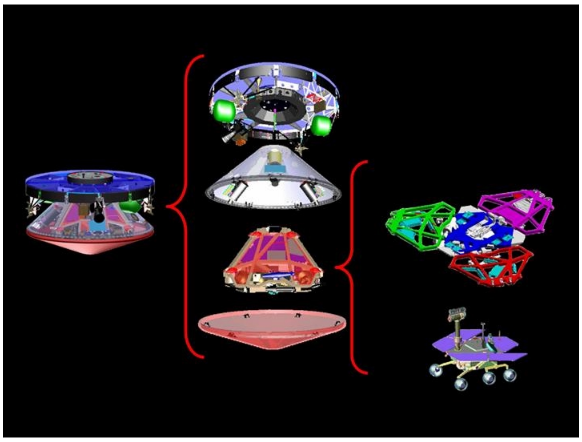

Fig. 5. The nested configuration of the rovers in flight to Mars.

The parallel team efforts during development continued into “cruise” (the period of time to fly from Earth to Mars) operations. Once the two rovers had launched, the cruise team was now engaged in flying two rovers to mars as well as participating in preparations for arrival at Mars. Many of the team experts could not be exclusively devoted to one segment of the mission and found themselves actually working all three.

The team focusing on Entry, Descent, and Landing (EDL) was nearing the end of the available time to prepare. The landing process was very autonomous- the spacecraft would jettison the cruise stage and then begin a six-minute descent to the surface of Mars from the time it entered the upper atmosphere. The spacecraft’s actions during those six minutes were entirely autonomous. Thus, the “critical sequence” for landing was tested over and over under a wide envelope of conditions (including a significant amount of off-nominal conditions).

As the two rovers were approaching Mars, the flight team began a series of detailed Operational Readiness Tests (training) to prepare the staff for landing day and for landed operations.

The MER mission was planning to work on “Mars Time” during some period of landed operations. The Martian day is approximately 39minutes longer than the Earth day. Thus, in order for operations teams to stay in synch with the times on Mars during which the Rover as awake, it was necessary for the operations teams to live on Mars Time just as the Rovers would be. As the teams began conducting operational readiness tests on this shifted time schedule, there were many rapid lessons learned about the extent to which other infrastructure would be necessary to support this effort such as food service, etc.

Just as the team had begun to learn how to fly two rovers while doing all the other development activities, a record-breaking solar storm occurred that was causing interference with many of the interplanetary spacecraft throughout the solar system. The only way to be sure that the rovers were protected from possible damage from the storm was not to command during the periods where the solar spots were in view. This cut the available commanding time in half that could be used to complete pre-arrival activities on both vehicles. Thus much of the operational timeline for completing the activities needed before landing had to be re-done to accommodate this new radiation-related commanding constraint.

5. LANDING & SURFACE OPERATIONS

The landing sites for the Rovers were chosen after years of discussions between scientists all over the world. Many of the prominent Mars experts had developed, over the years, favorite locations for potential sites on the planet that might have evidence of past water. The MER mission was an opportunity to test those theories. A series of workshops were conducted during the development of the Rovers to select among those possible landing sites.

The final landing sites that were chosen were a compromise between engineering safety and the locations of highest likelihood to reveal past water. Gusev Crater and Meridani Planum were to be the landing sites for Spirit and Opportunity respectively. Mars Global Surveyor had found evidence at Meridian Planum of hematite, a mineral known to frequently form on Earth in the presence of liquid-water. Thus, the location became a clear front runner for a landing site and a prime example of information from an orbiting spacecraft being used to guide the location of a landed mission.

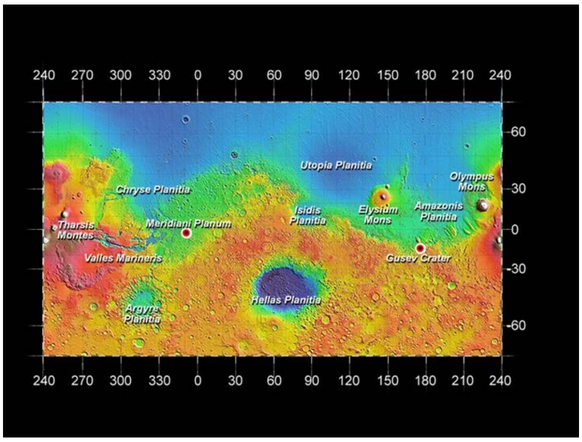

Fig. 6. The landing locations for Spirit and Opportunity.

The landings themselves were exercises in self-control-watching something over which one has no control play out hundreds of millions of miles away. In addition, Mars had been subject to a large-scale dust storm shortly before arrival that changed the density of the atmosphere- a critical environmental variable upon which many of the EDL parameters relied.

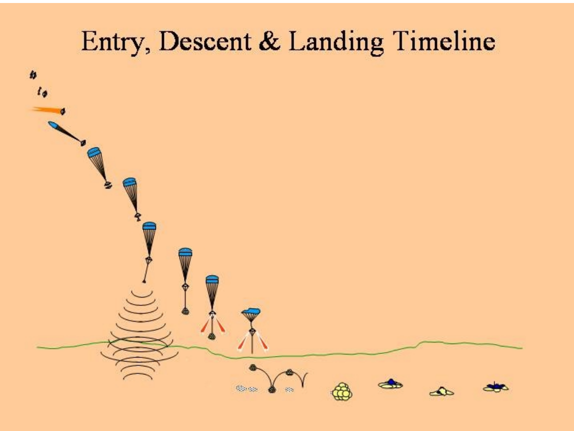

Fig. 7. The 6 minute descent to the surface

Spirit landed on the surface of Mars and suddenly the reality of Mars time hit the team. At this point, we had one team living on Mars time and the other not. The Martian day or "sol" is 39 minutes longer than the earth day and in order for the operations team to stay in sync with the rovers, we had to go into work 40 minutes later every day.

This became especially pronounced when Spirit suffered a serious anomaly and one team was involved in working around the clock to recover the rover and the other team was still operating on earth time landing the Opportunity rover

Once both rovers were on the ground and operating, the team really never saw each other due to time changes. Managerially, living on mars time was hard for the team and their environment had to be made as easy as possible-with everything from clocks to light shades to keep out as many earth-based distractions as possible.

Spirit was now exploring the rocky terrain at the Gusev landing site while Opportunity was exploring Meridiani Planum. In an amazing development, Opportunity had not only landed safely on Mars but had landed inside a crater with exposed bedrock only meters away.

Fig. 8. Opportunity Landing Site with exposed Bedrock

Exposed bedrock would be a key component to achieving the science goals of determining if past liquid water exists on Mars. The initial results from Opportunity were extremely encouraging and in March of 2004 the team announced the successful completion of the main mission objective-Opportunity had returned conclusive evidence that there is past evidence of liquid water on Mars.

There was also an extensive learning curve on now to control the rovers, Now that we knew the landing sites and terrains, the rover navigators could begin to customize the driving approach for the terrain, At first the rover planners drove only short distances while learning how to use auto-navigation. Meanwhile, the engineers at the Opportunity landing site were learning the details of fine driving while their compatriots at Gusev were having all the fun driving the long distances. The two landing sites were so different in terrain- the rocky nature of Gusev vs the flat plains of the Opportunity landing site- that the rover drivers were gaining invaluable experience in moving rovers around on Mars.

The Mars Pathfinder mission in 1997 broke ground for how to drive a rover on the surface of another planet but these rovers could drive in a day as far as the Pathfinder Rover drove in its whole mission.

In addition, these Rovers were not remotely controlled (a.k.a. a joystick) but rather remotely sequenced. The light-time delays in commanding the rovers as well as the intervals without ground contact meant the rovers needed to receive their instructions in the martian morning and then carry out the commanded sequences without intervention from the ground. This had clear implications for the surface driving.

The rovers have two driving modes- using hazard avoidance or having it turned off. Hazard avoidance uses the front and rear hazard avoidance cameras (HAZCAMs) to develop an elevation map that then allowed the navigation software to use to determine location of obstacles and to instruct the mobility system to "go around" the hazards. When the rover is driving in this mode, it clearly takes longer due to the need to develop these maps. If the rover has hazard avoidance turned off, it can drive more rapidly but has to rely on the ground operators knowledge of the terrain ahead and "drive blindly" trusting that the commanded sequence will not take it into an obstacle (there is also onboard fault protection that protects the Rover from harm).

Given the varied experiences the rover drivers were getting from the two landing sites, it was soon possible to consider tuning the rover navigation systems as the team became more experienced at driving the rovers. Indeed, there were a number of new flight software loads done to make improvements in the way the Rovers were driving.

There were a number of times during the surface operations phase where the control capabilities of the payload suite were stretched. One of the most far-reaching of these was using the Instrument Deployment Device and the Micro-Imager in an unanticipated way. The scientists were closing in proof that the past water at Meridiani Planum had not only been surface water but that it was actually standing water. One key clue that seemed to be emerging was the presence of "layering" in the rocks similar to what might be seen in a sedimentary layer on the earth at the bottom of a lake. In order to prove this, the scientists needed a full scale look at one of the rocks in the outcrop. However, none of the high – resolution cameras on board could provide such a close-up look at the rock and the Mico-Imager on the end of the robotic arm had to small a field of view. Thus, in a never-tested use of the robotic arm, a series of images were taken in a Micro-Imager mosaic that allowed a detailed image map of the rock to be made and provide the proof of layering the scientists needed.

7. EXTENDED MISSION OPERATIONS & SUMMARY

Both Rovers have survived well beyond their design lifetime of 90 days. Though both Spirit and Opportunity are starting to show signs of wear mechanically, the rover power profile is holding steady.

One Rover has proven to be an apt crater explorer while the other has become the first robotic mountain/hill climber on Mars as Spirit treks upward into the Columbia Hills.

As operations have continued well beyond the 3 month point, the teams have become smaller and more cross-trained. At the start of the mission, engineers had more focused roles in particular areas of operations but have gradually taken on broader areas of responsibility to account for the smaller staff. In addition, the science teams are now working almost exclusively from their home institutions- another process that had to be refined by trial and error.

The Rovers have both already survived a Martian winter and the team looks forward to continued operations.

8. SUMMARY

The Mars Rovers were a mission borne out of necessity- the desire to take advantage of an excellent landing opportunity in 2003 and the need to re-use a landing method done successfully in 1997. By building one a heritage system and then adding new capability using engineering teams with focused objectives, the Rovers launched on time. Since landing, the team has steadily improved in their ability to operate the rovers efficiently and new continued improvements in surface navigation have been critical in the rovers' success at finding evidence of past water. As with any interplanetary robotic mission, the lessons learned from these vehicles (both scientific and engineering) will feed forward into the upcoming Mars missions as well as other solar system exploration in search of water.

9. ACKNOWLEDGEMENTS & REFERENCES

The author gratefully thanks the entire Mars Rover Exploration Development and Flight Team whose work this paper chronicles. The author is privileged to have been a part of the team and to relate there story herein.

The research described here was carried out at the Jet Propulsion Laboratory, California Institute of Technology and our mission partners and was sponsored by the National Aeronautics and Space Administration.