Presented at the Design Theory Workshop of the International Design Society, Mines Paristech (School of Mines of Paris), 1-2 February 2010.

A Design Framework for Integrating Instruments, Software, Processes, and Organization in Telerobotic Mission Operations: Experience with the Mars Exploration Rovers

William J. Clancey

NASA Ames Research Center
& Florida Institute for Human and Machine Cognition
Intelligent Systems Division, MS 269-3
Moffett Field CA 94035
William.J.Clancey@NASA.gov

Introduction

Most scientific space exploration involves using instruments onboard spacecraft (e.g., Cassini orbiting Saturn, Hubble in Earth orbit) or surface vehicles (e.g., the Mars Exploration Rovers, Huygens on Titan). The design of such missions ("mission operations design") requires that systems engineering properly relate hardware (e.g., instrument "payloads"), operations processes (i.e., from receiving telemetry to transmitting commands), organizational roles, and software systems (e.g., configuring command sequences and tracking data from request to interpreted products). Usually, the instruments are "telerobotic"—computer-controlled and regularly reprogrammed by international teams of scientists and engineers. Configuring targets and parameter settings for the instruments constitutes another design process ("science operations design"), which occurs throughout the mission. These missions continue for years, so decisions made early in the mission operations design process affect the complexity, cost, and quality of science operations; spacecraft hardware designs particularly affect how scientific disciplines collaborate and whether they can use instruments in complementary ways. Mission failures serve as another motivation for improving design practices. A study of 35 mission failures[i] shows patterns in inadequate requirements analysis, use of legacy systems, and management supervision that suggest mission operations design could benefit from a holistic theoretical framework for designing complex systems. This paper introduces a design framework that represents the mission operations experience of the Mars Exploration Rovers. For contrast, Cassini's instrument mounting is presented as a less parsimonious design that complicates scientific collaboration.

The Mars Exploration Rovers

Planetary surface investigation with a mobile robotic laboratory (i.e., a "rover," Figures 1 and 2) requires tools for navigation as well as a mission operations process that allows recoordinating the remote system's behaviors on a daily basis. Jake Matijevic, the Engineering Team Chief for MER Operations at the Jet Propulsion Laboratory,[ii] who focused on systems engineering throughout the mission, explains the operations design constraints:

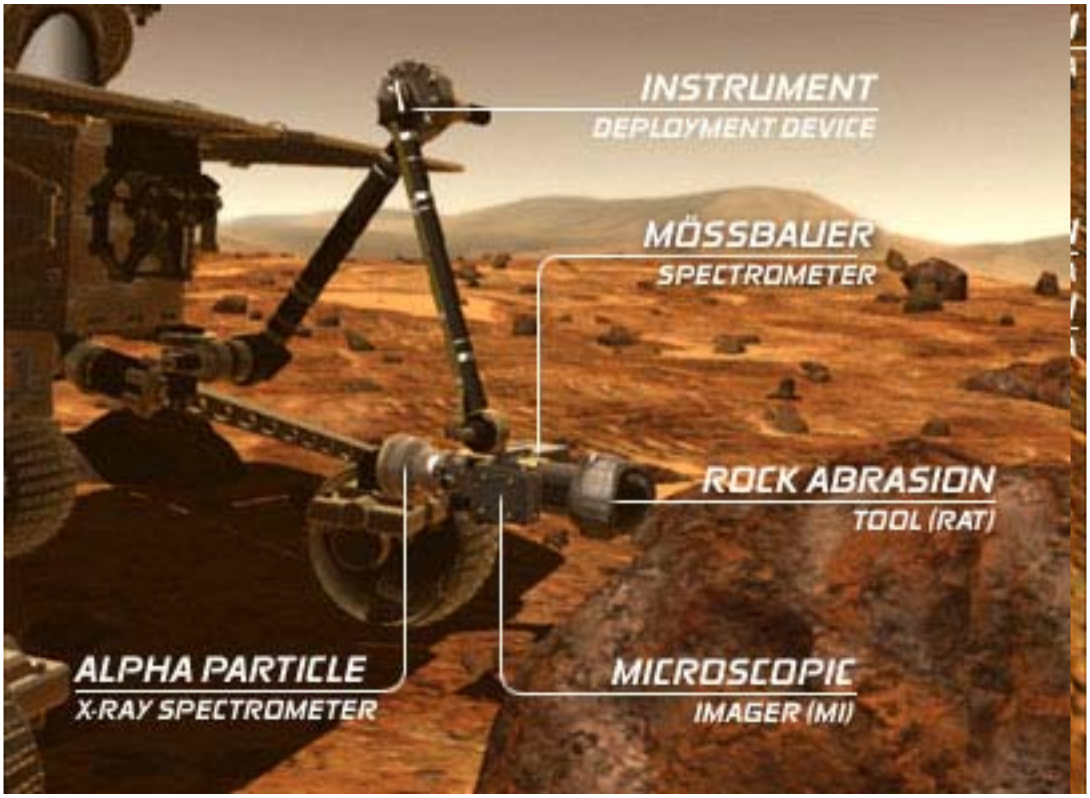

FIGURE 2. MER Instrument Deployment Device (IDD) and instruments. (Graphic: ,) NASA/JPL)

> When you are in the situation where you're trying to do these missions on a rover, you actually *require* information that is generally not paramount when you're flying a spectrometer or an imager or something around or by the planet. Here understanding the engineering that's possible with the device is essential for the scientists to acquire the information they are interested in. In other missions [the science operations plan] becomes part of the initial design and once the instruments and the vehicle are deployed, either as an orbiter or flyby, you're actually *clocking* through the mission at that stage. The things that you do with the spacecraft are kind of well-understood, even at the various design review stages, many years before the mission, before the instrument is ever deployed. In a rover mission, you are effectively creating "design behaviors." And now you're plopped down in an environment about which you have little if any information, and given the charter to go forward, do the exploration, find the opportunities to use your instruments. That requires a real understanding both of the engineers of the scientists and the scientists of the engineering: What are the capabilities and what is possible to be done?

The twin Mars Exploration Rover (MER) missions, called Spirit and Opportunity, each initially cost about $400 million and were in planning through a laborious set of proposals over nearly a decade. The principal investigator, Steve Squyres from Cornell University, brought the team together through a complex obstacle course of NASA programmatic changes, in large part caused by the twin failures of Mars orbiter and polar lander missions in 1999. The team, originally called Athena, and still so-named internally, often learned of each other's capabilities by competing for earlier canceled or revamped programs. They merged subteams with innovative instrument technologies, and incrementally developed a conception they called a "robotic geologist" that could enable them to perceive and act on Mars, which they characterized as being a "surrogate."[iii] The scientists wanted an exploration system that would enable them "to do real field geology on Mars."[iv]

The MER "exploration system" comprises instruments, software, processes, and an organization that enables the team to do planetary field science by operating the mobile laboratory remotely from Earth. For example, describing the origin of the RAT (Figure 2), Steve Squyres (Principal Investigator [PI] of MER) said, "Our rover was supposed to be a robot field geologist. When you see field geologists on Earth, they've got their boots, they've got their backpacks, and always, they've got big rock hammers."[v] Correspondingly, the microimager (MI) is like a powerful hand lens, the wheels can be programmed to dig trenches (like scraping your boot in the dirt), a brush on the RAT can sweep away dust, Navcam cameras provide a ground view, and the Pancam on a mast provides stereo images at about 1.5 m (Figures 1 and 2). The RAT, brush, and MI are mounted on an arm with an elbow. Its reach is like a person's. Using the arm is like kneeling, the mast enables standing up and looking around—though weight and packaging constraints forced it to be less than the person-height originally desired.

Crucially, a field geologist's set of tools need to be designed so they can operated in a coordinated manner, allowing data to be correlated. For example, MER's filtered stereo images can be usefully related to its molecular and atomic scans. Areas for geochemical analysis are first scraped to form a smooth surface and imaged in detail. This design followed from experience on Viking and Pathfinder, in which analyses were biased by the dust; the RAT enables probing to find the composition of the rocks.[vi]

Despite these analogies between the rover and a geologist's tools, three additional, highly sophisticated instruments for detecting iron compounds and analyzing chemical element composition—Mössbauer, Mini-TES, and APXS—make each MER more like a *robotic laboratory*. Geologists do not take such instruments into the field, but rather analyze their rock and soil samples in laboratories at home. Using such tools requires hours of "integration" for the sensors to "read" landscapes and materials on different scales. For example, Mini-TES pixelates a frame of view by individually programmed readings, moving a mirror to make multiple scans of an area, typically requiring five to 30 minutes, but sometimes more than an hour. An APXS integration could require 10 hours. Such programmed analysis may require considerable power and time, key constraints in daily planning.

Taking time to configure, acquire, and evaluate data from the laboratory instruments frustrated field geologists who were more familiar with broad surveys before considering geochemical analysis. Thus, Spirit's team took 156 sols (a sol is a martian day, about 24 hours 40 minutes) to reach the base of the Columbia Hills, 3.25 km from its landing site, and then another 34 sols to reach its 100 m high summit—which one field geologist said he would have climbed on the first day. (Indeed, the original planning for MER assumed a 90 sol nominal mission with a 600 m combined distance for the two rovers, making the Columbia Hills seem initially too far to reach. The results illustrate how MER's engineering design was based on minimal requirements and the unpredictability of how systems will perform in a new extreme environment.) When MER scientists say they "could have done everything in a day" they are referring to the traverses, not the scientific data collection and analysis. This inherent conflict between the methods of laboratory and field scientists illustrates part of the challenge in crafting the exploration system: Diverse scientists were forced to ride together on the rover for five years, deciding as a team where to stop, what instruments to apply, and when to move on, and thus forced into the anonymity of a communal experience.

Design Strategy of the Science Systems Engineer

From the beginning in the early 1990s in conceiving instruments and rover proposals, Squyres saw his job not merely as a scientist, but as a *science systems engineer*, someone who designed and tested the rover holistically, connecting parts and processes to ensure the quality of the science:

> I saw my job during development as kind of science systems engineering—to be able to look at the total system and find the parts of the system whose quality one way or another has the greatest bearing, the greatest influence on the science. Because what would happen is that the engineers would each be given part of the problem to solve, and they would go off and solve their part of the problem. But they didn't maybe have the broad overview of what we were trying to do scientifically. They had a specific job, a set of requirements to meet, to build a bracket, build a widget, build a wheel, build a computer, build a rover that would meet those requirements. But, they might not have enough insight into the science to realize that even though this one part of the rover meets the requirement, if you could make it just five percent better or ten percent better it would have a huge influence on the quality of the science. And there were some things that were tremendously important and others that were much less so.

Another example of high leverage is the speed of image processing on board the rovers:

> We didn't do a good job when we first started defining the requirements of talking about how *fast* the software for doing image processing on board the rovers had to be, and initially it was painfully slow. If you wanted to take a series of Pancam images you could only take one image every six minutes. We got that down to thirty or forty-five seconds. If you want to take these great big panoramas and you have only so many minutes during the day, that's an order of magnitude difference in how much data you can acquire. And so I screamed bloody murder about that one forever.

In his role as science systems engineer, Squyres had responsibility for twelve different payload elements, six instruments on each rover, which he shepherded through the design, test, verification, delivery and integration process.

Strikingly with Squyres as the sole PI in the MER project, the design of the science instruments and operations could be under his control—his job description said, "wholly and solely accountable for the Athena science payload instruments"—a freedom and responsibility offered by NASA's Discovery program:[vii]

> Our mission was structured very differently from most. The norm on planetary missions—Cassini, MSL [Mars Science Laboratory]—has been that you have individual PIs for every single instrument. And then they compete with one another for resources. You have a project science group that is chaired by the Project Scientist.[viii] And they have the job of herding these cats—it's more like herding saber tooth tigers I think—getting everybody to play nice together, to work together. The situation on M-E-R was fundamentally different in that we had a single PI. And so I had responsibility for the entire science payload.

The consequences for the quality of the overall design were sharp and fundamental:

> It meant that I could optimize the performance of the total science payload. It was a science systems engineering task where you have this very complicated multidimensional problem. But I had a lot more levers that I could push as PI for the whole payload, than a typical PI does for a single instrument. So for example, if I had one instrument that's doing well and one that's in trouble, say financially, I can move resources from one to another. I can push money around, I can push people around, I can do the things that I need to do to solve one problem. It's a system optimization problem, and the higher the level at which you do the systems optimization, the more parameters you have at your disposal to play with.

Crucially, Squyres' oversight and control focused on the integration of the instruments—so rather than worrying about how independent PIs worked together to produce a coherent result, he could directly focus on designing the instruments to work together:

> When it came to the *design* of the payload, rather than have a bunch of PIs each go off, and they each design their instrument and then they propose it, and you've got to figure out some way to make it work together, I was able, at the *outset* to design the microscopic imager, the RAT, the Mössbauer, the APXS, so they had fields of view that fit together nicely so that everything worked when you went to look at a given spot on the Martian surface. Just design the whole system so the pieces fit together. And that was nice!

This idea of integrated instruments has deep implications throughout the project. It means not only that the instruments align in useful ways, but it extends to the choice of instruments, so the rover as a whole is designed to remotely carry out the scientists' work. Correspondingly, the concept of the rover as a geologist's physical surrogate has ripples throughout the design and experience of the mission: 1) the scientists are organized according to disciplines not instrument payloads; 2) the scientists as individuals can project themselves into the rover's body, to orient themselves to the work setting and imagine what they would do next; 3) the metaphor of the "robotic geologist" shaped stories about the mission, creating a kind of romance about the scientists' expedition, and this personification of the mission reinforced the science team's identity and focus on the common goal of understanding the planet Mars.

"One Instrument, One Team"

Squyres started with a simple, logical concept of the rover as being a single, coherently operating entity, hence the phrase "robotic geologist." This concept plays out in the selection of the instruments, the systems engineering to make sure their operations are complementary, and the visualization tools. The holistic concept of the rover flows naturally into the design of the organization for collaboration—summarized by the rubric, "one instrument, one team":

> It was just intuitively obvious to me given the payload that we had put together. You've got these sensors and each of them provides complementary bits of knowledge, so that the totality is more than the sum of the individual parts. *You're going to use the payload to fullest advantage, if people look at it as being entirely at their disposal.* So what happens is that I'll have geologists or geochemists come in to the meeting and say, "Well, we can really understand this if we first take a Pancam of it, hit it with Mini-TES, and if it looks like this, go over and APXS it." You know, that's the idea! The whole idea behind MER is that these tools work together. Look at the silica discovery. Okay, the mobility system, which we use as a soil physical processes tool, trenches up some soil. We notice it with Pancam, we hit with Mini-TES; it looks interesting, and we go over and we figure out what it's made of with APXS. Everything works together.

Indeed, MER fits together on many levels in several ways: Sensors fit sensors, the rover fits a field geologist, disciplinary contributions fit each other. So rather than "my instrument," did Squyres intend for the scientists' mental model to be "my shared rover"?

> Exactly! And so, it was just really evident to me from the start that if you structured the team along *those* lines that you have geologists arguing with mineralogists arguing with atmospheric scientists. Well, that's what scientists do, naturally. If you were out there in the field, doing geology with you and your field partner, you might be arguing about what this rock means or what that rock means, but you're not going to argue about, "Well, should we use the rock hammer or should we use the Brunton compass?"

But given the originality of this concept and its variation from the instrument-centric organization of other missions, didn't he get any pushback? Squyres says, "I'm not sure that they even noticed that I pulled it off on them. *(laughs)* No, I didn't get any pushback. And it's worked." This unification of interests helped people share the rover for different interests, avoiding the plight of individual instrument teams that feel they are being slighted: "You use the instruments when it's time to use the instruments, in the service of the science you're trying to do. It's a pretty fundamental point about how the mission's set up."

This transparent discipline of "one instrument, one team" has a special manifestation in the decision-making process itself, as the investigation must also be a team effort, requiring actions to be justified to each other and systematized to endure scrutiny of many other scientists for years to come. This is another paradox of being a scientific explorer on Mars working remotely through a rover: The scientists had to work as if they were writing a textbook in the field, rationalizing instrument application through explicitly articulated hypotheses and investigation strategies.[ix] The discipline of this decision-making process was unexpectedly facilitated by the independent mission operations design decision to enable daily commanding: Commands were transmitted every sol, with sufficient data returned and analyzed to enable retargeting and programming the instruments by the next morning on Mars (i.e., "one sol turnaround").

Unexpected Benefits from Daily Commanding

Sending commands to each of the MERs every sol was originally viewed by the scientists as necessary to maximize the amount of scientific data gathered and exercise the rover's capabilities within the expected short lifetime. But a fundamental serendipitous benefit emerged: The intensive daily activity, with sufficient feedback to assess the previous day's plan and to move on, significantly increased the scientists' engagement with their rover's activities, making the rover's state and what might be done next something they thought about every day, month after month, and ultimately year after year. Simultaneously, collaboration was enhanced.

Staying engaged and projecting themselves into the landscape fits the scientists' experience of being in the field—a turnaround of more than two weeks as on the 1970s Viking Mars lander would instead require additional effort to remember where they were or what they were doing. Even with distributed operations, more than a year after the nominal mission had completed, Jim Bell (lead Pancam scientist) said that daily turnaround "helps keep us all excited about the exploration aspect of these missions: many times, we don't know what we're going to be doing that day with a rover until we see the results from the day before."[x] This routine daily process, operating a wide variety of instruments on a mobile vehicle, with people distributed across the world, continuing for years—inconceivable during Viking's time and only prototyped during the 1997 Mars Pathfinder/Sojourner mission—had been made possible by the advances in constraint-based planning and scheduling programs, shared databases, and web-based graphical browsers developed over the preceding decade and refined in training as well as during the mission.

In short, daily commanding improved the quality of the scientists' work, increasing their productivity in selecting suitable targets, which on a mobile laboratory has pronounced tactical and strategic aspects. Long-range planning for MER has meant scanning for or anticipating features not in the immediate landscape and has therefore been more scientifically conceptual and strategic than on the Viking lander (in which all plans were 16-20 sols out, involving features already in view). Using a rover requires projecting time and resources into unknown places, but daily commanding allows a luxury of separating daily operations from thinking about the scientific investigation of the topography.

Balancing disciplinary interests of the scientists, promoting fairness and hence lowering competitive stress, was also facilitated by daily commanding. With the handful of instruments available and the competing interests of scanning the atmosphere, doing detailed geochemistry analyses, and moving on in a broad geological survey, being able to tradeoff opportunities every few sols enabled the process to be cooperative and enhanced the experience of working together as a multidisciplinary team. Consequently, the daily command turnaround helped implement Squyres' vision that the MER team would not just cooperate in sharing a common platform, as in other missions, but would collaborate in sharing a common inquiry, namely studying Mars holistically.

In conclusion, maximizing productivity, planning as field scientists, maintaining awareness and engagement (virtual presence), and making the use of the rover a coherent single scientific investigation were all enabled or enhanced by daily commanding. This is an important lesson learned because aside from productivity, *none* of these benefits were listed in the formal presentations to justify daily commanding during the MER design review meetings in 2001.

A Design Framework for Scientific Exploration Systems

Figure 3 summarizes the design framework of the MER exploration system. At the top we place the primary measurable objective of the mission, the quality of scientific work. The scientific work has two opposing aspects: *systematicity* (required by the laboratory scientists, e.g., regular sampling on the traverse to the Columbia Hills) and *opportunism* (required by the field scientists, who want to examine outcrops and major features seen from a distance, e.g., the Block Island meteorite spotted during Opportunity's traverse on the Meridiani Planum during August 2009).

Being systematic is supported by a deliberative, textbook approach and by the systems engineering of integrated field and laboratory instruments that enable thorough sampling (e.g., when studying individual rocks and outcrops) and periodic sampling (e.g., characterizing geochemical variations along a traverse, including along or down a crater rim).

Being opportunistic is enabled by knowing what scientific operations can be performed on Mars (*situation awareness*) and timely execution through commanding tools (software for planning and programming instrument operations). Situation awareness in turn requires long-term participation by individuals (*continuing engagement*) as well as oriented perceptual-motor projection into the immediate setting to engage field investigation and laboratory sampling skills (*virtual presence*).

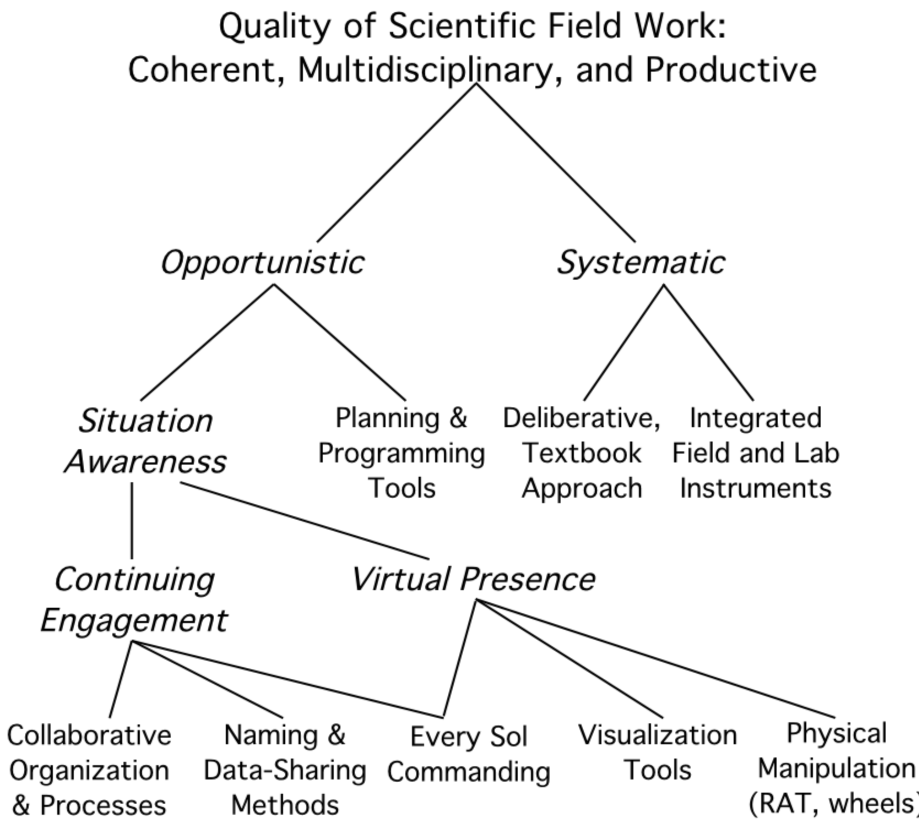

FIGURE 3. Overarching requirement for coherent, multidisciplinary and maximally productive field science is related to *design features* of the MER exploration system (italics), supported by the more specific methods employed.

Many more details could be included in this summary of MER's mission design, such as having multiple communication opportunities during the sol, daily retargeted imagery, visual alignment of images and measurements with multiple perspectives, documentation of rationales for observations, and integrated software tools.

Surprisingly, MER's design exemplifies how tools that promote collaboration might have little or nothing to do with communication media. Conceiving the rover holistically, according to the "one instrument, one team" mantra (in contrast with a collection of instruments that shares a common platform), promotes multidisciplinary collaboration for performing field science on Mars, forging a relationships among field scientists (who may have never used certain instruments in the field or performed such laboratory analyses) and laboratory scientists (who may have never participated in a field expedition). The following brief case study of the Cassini mission illustrates how a hardware design not as well integrated with scientific operations makes collaboration among disciplines more difficult.

Cassini: Example of a Hardware Design Lacking Parsimony for Operations

Many people (including space scientists and engineers) might not understand, even after years of operating Spirit and Opportunity, how fundamentally different a rover mission is from operating an interplanetary spacecraft. As we have seen, the MER team engages in a collaborative process of selecting and studying well-targeted locations by moving all of the instruments together from place to place. On Cassini, a spacecraft currently orbiting Saturn, different discipline teams study the planet, its rings, magnetosphere, Titan, and icy satellites.[xi] These teams effectively own the operation of the spacecraft at different times and places in the orbits of Saturn.

Matrixed with the discipline teams, Cassini is organized into twelve *science instrument teams* conducting distributed operations across the United States and Europe, each of which includes international engineers and scientists who participate in operations.[xii] (The Huygens Titan probe, with six more instruments, was managed as a separate ESA project.) Unlike MER's single PI, Cassini has 12!

> Each principal investigator (PI) is solely responsible for the design, construction, integration and flight operation of his or her own instrument, including mission planning, sequencing, instrument monitoring and science data acquisition, processing and analysis. Uplink and downlink processes have been developed expressly to support the distributed nature of the project.[xiii]

In carrying out a "remote sensing" investigation during a nominal 4 year mission including 75 orbits of Saturn, Cassini's payload was conceived as a *collection of instruments* that independently collect a variety of data. In contrast, the MER scientists are investigating targets together on a surface landscape in a cross-disciplinary way—as with Cassini, they must agree where the platform will be pointed and whose instruments will be operating, but they do so by going places and doing things together through a physical surrogate conceived as being a single entity. The MER team operates by consensus at the level of what is being investigated and how the instruments are applied to resolve scientific questions, hence their organization into science theme groups. Managing shared resources and the engineering, pragmatic level of operating the rover are resolved with respect to what's best for the shared inquiry, rather than allocated to individual PIs.

Comparing a planetary flyby or orbital mission to a rover mission like MER involves making a distinction between *cooperation*—sharing common resources (e.g., power, time, bandwidth, the payload platform)—and *collaboration*—engaging in a common inquiry from different scientific perspectives. Generally speaking, Cassini's PIs are cooperating, not collaborating. Reflecting the topographic regions of the Saturn system, Cassini's discipline groups share the spacecraft resources and time by partitioning orbit segments, and then engage in "collective planning."[xiv] Thus scientific work has been partitioned by hardware and the region being studied, rather than basic scientific knowledge (geochemistry, geology, mineralogy, meteorology) as on MER.

Cassini's resource-coordinated, independent inquiry occurs on MER, too. For example MER atmosphere studies are largely independent of geochemistry studies. But MER's instrument teams are more often collaborating in deciding what features to investigate. Grounded in a common geography, they are exploring Mars together.

The wildly different kind of phenomena being studied during Cassini's mission—storms on Saturn, gravitational perturbations in the rings, methane rain and erosion on Titan, ice geysers on Enceladus—suggest that the nature of a phenomenon may dominate scientific discipline in mission operations design. The design of another interplanetary orbiter supports this conclusion: Galileo, a Jupiter orbiter (1995-2003), had seven PIs for the pod of atmospheric probe instruments, six PIs for the six fields-and-particles instruments, and four PIs for the remote sensing instruments.[xv] And so we can generalize the "robotic geologist" metaphor to include the lessons from Cassini and Galileo: *The topography and regions being explored may dominate the design of the instruments, determining what forms of joint activity are desirable and possible.*

In short, it appears plausible that how science operations are partitioned on Cassini is inherently driven by remote sensing of different planetary bodies (imposing specialization by instrument and regional interest) and the joint NASA-ESA mission (imposing a distribution of scientists living across ten time zones). However, a hardware design decision lacked parsimony with respect to operations, making cooperative planning more difficult than it needed to be, given available technology:

> The spacecraft was originally planned to be the second three-axis stabilized, RTG-powered Mariner Mark II, a class of spacecraft developed for missions beyond the orbit of Mars. *Cassini* was being developed together with the *Comet Rendezvous Asteroid Flyby* (*CRAF*) spacecraft, but various budget cuts and rescopings of the project forced NASA to terminate *CRAF* development in order to save *Cassini*. As a result, the *Cassini* spacecraft became a more specialized design, canceling the implementation of the Mariner Mark II series.[xvi]

Because of the budget constraints, the hardware design team decided to mount the instruments directly on the body of the spacecraft ("body fixed"), rather than providing the original two articulated scan platforms, which by the Mariner Mark II design could have been directed independently of Cassini's antennas. The articulated, dedicated Huygens antenna was also deleted.[xvii] A NASA-funded analysis team studying Cassini's design concluded that "the spacecraft was not designed for maximum operability."[xviii]

Without a scan platform, the engineers must turn the body of the entire spacecraft to bring a target into an instrument's field of view (however, several instrument teams added their own articulation), and thus the scientists can generally can use only one instrument at a time, leading to a complex trade system. This method also uses far more fuel for maneuvering the spacecraft, shortening the duration of the mission, and thus reducing scientific return. Furthermore, the instruments have "significant power limitations resulting in the need for operational modes that manage the total power needs of the entire flight system including the instruments."[xix]

Perhaps unavoidably, the Cassini scientists' engagement in the mission is at the opposite end of the spectrum from MER's daily contingent planning. Cassini's trajectories may be planned up to a year in advance because they are constrained by orbital physics and the positions of the moons—making route planning very different from an overland expedition like MER. However, on the plus side, favoring cooperative planning, Cassini plutonium RTG power system enables it to be operated at any time (rather than just when the sun is shining on solar panels), and the nominal mission was four years. Some opportunistic retargeting is also possible, for example flying through the icy plumes of Enceladus multiple times in 2008. However, the long waits between requests and getting data for Cassini—imposed by physics, the geographically defined investigations, and the spacecraft's design—engender less sense of working together as a single "Cassini Project Team" than for MER, requiring the patience of waiting in turn (besides waiting to get somewhere), and requiring more effort to reach fair operations plans.

In conclusion, the Cassini scientists were inherently isolated by their interests in different phenomenon located in very different locations as well as physically working in different times and places on Earth. We can speculate that the teams would have been more likely to collaborate (that is, work on each others' personal projects) if the original scan platforms had been provided. Teams might have had more flexibility to plan and design complementary science operations, allowing more data to be collected from different instruments during optimal sensing opportunities near Saturn, a moon, or the rings. Whether the undesirable spacecraft design has affected the quality of contributions to planetary science remains to be seen. Tradeoffs are difficult to assess because adaptations are possible, and a mission can be successful despite lacking optimal relations among hardware, software, processes, and organization. An observational study of the science team is necessary (and underway[xx]) to examine how the team realizes the ideals of the scientific method within the pragmatic constraints of the mission. Simply judged by photographs that reveal the physical mysteries of the Saturn system, the results from Cassini so far are stellar.

Conclusions

Work system design involves properly relating people and machines—designing hardware, software tools, organizations, and processes holistically, so influences are complementary, rather than components that people must continuously struggle to fit together and workaround.

MER and Cassini are ongoing highly successful missions. The two case studies illustrate how space science operations expose both helpful and antagonistic aspects of a mission operations design. Hardware decisions cannot be changed after launch and will affect a mission for perhaps 5 to 20 years. Software and spacecraft settings are changeable, but reprogramming is prone to errors that only show during operations and are sometimes fatal (e.g., Mars Polar Lander). Yet the most important aspect, the quality of the scientific work depends not just on survival of the systems and remote control to acquire data, but collaboration among multiple disciplines whose interests vary across interplanetary regions (e.g., planets vs. moons), fundamental analytic topics (e.g., geochemistry vs. climatology), and sensing instruments (e.g., spectral vs. micro-optical). This collaboration is not just a momentary interaction between people, but an ongoing state of mind, an attitude of commitment and a manner of participating, a way of conceiving roles and disciplinary relationships—which we found in both MER and Cassini can be reinforced or hindered by hardware decisions.

Repeated near-critical mission failures (e.g., MER/Spirit's memory problem shortly after landing in 2004; Phoenix's short-circuit in the TEGA oven door in June 2008; LCROSS's unnecessary thruster firing that used half its fuel trying to lock onto a star in September 2009) and cost overruns (e.g., Mars Science Laboratory consuming and delaying the budget of many other missions) suggest that mission operations design is an immature and unreliable process. The emerging field of "design theory" within a variety of engineering disciplines could lead to better processes and tools for managing complex interactions throughout the mission life cycle.

The strong visible relations between systems and people in the science operations of MER and Cassini (on top of NASA's compendium of lessons learned from failures) raise the possibility of developing a design theory for mission operations. We begin with the caution that a work systems design, in contrast with an artifact viewed in isolation, does not deterministically control the designed system's operation. In particular, procedures can be prescribed, but how people behave in practice cannot be fully specified or tested in advance: The work system adapts during operation through human learning and improvisation, as well as redesign of tools, organization, and protocols during the mission. Crucially, the tools in situ (e.g., on a planetary surface) interact with an unknown environment, with unpredictable results (e.g., LCROSS's failure to send up a visible plume on impacting the moon in October 2009). Scientists may also make discoveries and opportunistically change requirements during the mission (e.g., engineers had to learn how to steer MER on crater slopes). In effect, scientists' interpretations and plans, as well as the unanticipated (indeed, by design *unknown*) space environments spacecraft and rovers encounter are contextual aspects that make a priori evaluation of designs unpredictable. Rather, what is required are design principles that respect and positively reinforce known *relationships* among people, hardware and software systems, and the environment being explored.

Referring to the themes proposed for the Special Interest Group on Design Theory of the International Design Society,[xxi] a design theory for mission operations might be developed along the following dimensions:

- **Design theory and new approaches of high level and flexible structures of knowledge:** With respect to the study of "basic structures of reasoning and cognition," creativity, problem solving, and collaboration should be related to the social aspects of science operations, examining particularly how hardware (e.g., instrument and how they are physically controlled and aligned) and software tools affect multidisciplinary investigations and theorizing.
- **Theory-oriented laboratory experiments:** Regarding the field of Design Research, every space science mission provides a major opportunity to study how designs play out and are adapted in practice. Design researchers, social scientists, and historians should always be funded to participate in science operations as part of the team. Studying how scientists and engineers cope with fixed constraints and limited resources in configuring instruments for making observations—either collaboratively or cooperatively—is a wonderful opportunity to increase our knowledge of design practice. In effect, mission operations provides a natural laboratory, in which scientists and engineers (often distributed) are creating and experimenting with designs on a regular basis, usually for years.
- **Increasing the innovative and creative capacity of design methods:** Finally, we have learned that tools for designing science operations (e.g., what the MER rover should do tomorrow) inherently must incorporate engineering design methods. Scientific exploration in space involves programming a remote laboratory in space or a planetary surface. Scientists' flexibility and creativity is greatly enhanced by experiment design, operations planning, and

instrument configuration tools that automatically and transparently relate system constraints across levels of operation (e.g., power, memory, bandwidth, timing, and component dependencies). Such tools have advanced significantly over the past decade—visualization tools prototyped for Pathfinder in 1997 and planning tools prototyped for MER in 2004 were both routinely used by Phoenix scientists in 2008. Although their purposes are often prosaically described as "planning and scheduling," such tools are actually promoting scientific creativity—they enable engineers and scientists to collaboratively design a complex system of interactions among hardware, software, and organizational processes as they investigate distant worlds through robotically mediated sensing and acting.

In conclusion, space mission operations design is a field in which engineering serves the needs of science, balancing scientists' aggressive push into extreme environments and engineers' state-of-the-art spacecraft and instrument systems against the mutual requirements for survival and cost control. Planetary science has rapidly advanced from photographs to remotely controlled probes (e.g., Huygens) and wet chemistry analyses, though at considerable cost and with highly visible failures. Developing yet more sophisticated mobile laboratories, programming them efficiently, and investigating the solar system collaboratively and holistically will require a more systematic approach that relates academic research to the realities of packaging, testing, and operating complex systems in space.

Acknowledgements

This paper is adapted from a chapter in *Voyages of Scientific Discovery with the Mars Exploration Rovers*, submitted as a NASA Special Publication, with partial funding from NASA's ROSES History of the Science of Space and Earth Exploration program. I am grateful to the assistance of the scientists and engineers who participated in my study of MER's design and operations and to my colleagues in the MER Human-Centered Computing Research group who worked on the mission from 2001-2004. I thank also Jack Hansen, Ade Mabogunje, Mike Shafto, and Janet Vertesi for their comments and encouragement.

References and Notes

[i] "APPEL-Space Systems Development: Lessons Learned," *http://www.nasa.gov/offices/oce/appel/curriculum/courses/appel_ssd.html* (accessed 7 October 2009).

[ii] JPL MER Mission, "Baby Boomers on Mars," *http://marsrovers.jpl.nasa.gov/spotlight/20060320.html* (accessed 1 December 2007).

[iii] S. Squyres, *Roving Mars: Spirit, Opportunity, and the Exploration of the Red Planet* (New York: Hyperion, 2005).

[iv] Steve Squyres, quoted by A. Chaikin, *A Passion for Mars: Intrepid Explorers of the Red Planet* (New York: Abrams, 2008), p. 244.

[v] Squyres, *Roving Mars*, p. 81.

[vi] D. Dyar, "An Out-of-this-World Interview with Darby Dyar" *College Street Journal*, 12 December 2002, *http://www.mtholyoke.edu/offices/comm/csj/121302/dyar.shtml* (accessed 25 June 2009).

vii Squyres, *Roving Mars*, p. 31.

viii Joy Crisp was the MER Project Scientist. Her role was to manage NASA's MER Project, serving as liaison for the scientists to the larger JPL organization and NASA Headquarters, as well as representing the project for public outreach.

ix For elaboration, see Clancey, *Voyages of Scientific Discovery with the Mars Exploration Rovers*, Chapter 5, "Being a Textbook Scientist."

x J. Bell, "Photographing Mars." *Planetary Report* XXVI, no. 6 (November/December 2006): 12-18, p. 13.

xi T. L. Rey, "Project Science Group (PSG) Meeting #46 (October 31, 2008)," NASA JPL, Cassini-Equinox Mission *http://saturn.jpl.nasa.gov/news/cassiniinsider/insider20081031/* (accessed 11 October 2009).

xii S. Linick, C. Boyles, and P. Woncik, "Cassini Distributed Instrument Operations – What We've Learned Since Saturn Orbit Insertion," SpaceOps 2006 Conference. AIAA 2006-5820.

xiii P. Lock and M. Sarrel, "Distributed operations for the Cassin-Huygens Mission," *Space Ops 1998*, *http://hdl.handle.net/2014/22577* (accessed 11 October 2009).

xiv J. A. Vertesi, *Seeing Like a Rover: Images in Interaction on the Mars Exploration Mission* (Cornell University Science and Technology Studies Department, Doctoral Dissertation, May 2009), p. 171.

xv M. Meltzer, *Mission to Jupiter: A History of the Galileo Project* (Washington, DC, National Aeronautics and Space Administration Special Publication-4231, 2007), pp. 122, 131-132, 142.

xvi "Cassini-Huygens," *Wikipedia*, *http://en.wikipedia.org/wiki/Cassini–Huygens* (accessed 13 October 2009).

xvii J.-P. Lebreton and D.L. Matson, "The Huygens Probe: Science, Payload and Mission Overview," ESA Bulletin 92 (1997) http://www.esa.int/esapub/bulletin/bullet92/b92lebre.htm (accessed 11 October 2009).

xviii M. Holdridge, N. Pinkine, D. Artis, K. Hibbard, R. Lock, G. Tan-Wang, G. Welz, B. Jai, M. Shafto "Outer planets operations lessons learned," *AIAA Space 2009* (Pasadena, CA, September 2009), No. 6417, quoting p. 11. *http://pdf.aiaa.org/preview/CDReadyMSPACE09_2074/PV2009_6417.pdf* (accessed 26 October 2009).

xix J. Gunn, T. Anderson, and T. Sesplaukis, "Doing development right: Lessons learned from the development of the Cassini Saturn Mission Operations System," *SpaceOps 2004* (Montreal, Canada, May 17, 2004) *http://trs-new.jpl.nasa.gov/dspace/handle/2014/37850* (accessed 11 October 2009).

xx J. Vertesi, and P. Dourish, "The Social Life of Spacecraft," Laboratory for Ubiquitous Computing and Interaction, Department of Informatics, University of California, Irvine, LUCI-2008-09.

xxi SIG Design Theory, *http://www.cgs.ensmp.fr/DesignTheorySIG/* (accessed 15 September 2009).