A Retrospective Snapshot of the Planning Processes in MER Operations After 5 Years

Anthony Barrett, Deborah Bass, Sharon Laubach, and Andrew Mishkin

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive M/S 301-260
{firstname.lastname}@jpl.nasa.gov

Abstract

Over the past five years on Mars, the Mars Exploration Rovers have traveled over 20 kilometers to climb tall hills and descend into craters. Over that period the operations process has continued to evolve as deep introspection by the MER uplink team suggests streamlining improvements and lessons learned adds complexity to handle new problems. As such, the operations process and its supporting tools automate enough of the drudgery to circumvent staff burnout issues while being nimble and safe enough to let an operations staff respond to the problems that Mars throws up on a daily basis. This paper describes the currently used integrated set of planning processes that support rover operations on a daily basis after five years of evolution.

Introduction

On January 3 & 24, 2004 Spirit and Opportunity, the twin Mars Exploration Rovers (MER), landed on opposite sides of Mars at Gusev Crater and Meridiani Planum respectively. Each rover was originally expected to last 90 Sols (Martian days) due to dust accumulation on the solar panels, but Mars rapidly proved to be more benign than expected as gusts of wind periodically blew the dust off. Five years later, both rovers continue to explore the Martian surface. This paper is on the planning and scheduling tools currently used in MER operations to reduce effort and avoid the staff burnout that would otherwise result from the need to produce uplink sequences within a 10 hour deadline. How these tools are used and how they interact to support the operations process is the result of years of painstaking work by the MER Operations staff as the rovers have become better understood and their mission has evolved.

As illustrated in Figure 1, each rover has six wheels to traverse the Martian surface, a mast with three remote science instruments, and an instrument deployment device (IDD) – an arm ending with four local science instruments.

Copyright © 2009, California Institute of Technology. Government Sponsorship Acknowledged.

Figure 1. Mars Exploration Rover

For communications, each rover has a high gain antenna for receiving instructions from Earth, and a low gain antenna for transmitting data to the Odyssey or Mars Reconnaissance orbiters with subsequent relay to Earth. Given that it takes between 6 and 44 minutes for a signal to travel to and from the rovers, simple joystick operations are not feasible. Also, the energy needed for a rover to send data directly back to Earth strongly motivates relaying through the orbiters, which is only possible in fifteen minute windows that only happen four times a sol.

Initially, the assumption of a 90-sol mission resulted in the entire MER team co-locating at JPL and operating on Mars time (Mishkin et al., 2006). This approach involved a rover sending data to Earth late in its afternoon and then going to sleep. As the rover slept through the Martian night an operations team would analyze the data and devise the instructions for the next day. The rover would then receive these instructions upon waking up in the morning, perform a sol's worth of instructions, and then send the resultant data back to Earth to repeat this operations cycle on the next sol. The reason for this tight schedule came from the inherent execution uncertainty when dealing with unknown surfaces on Mars, e.g. a drive to a location might fail due to unexpected interactions when driving over sand covered in dust. Mars has a slightly slower planetary rotation than Earth, and this operations approach resulted

in people waking up approximately 40 minutes later every day, which was not sustainable for a longer mission.

To facilitate a longer mission, two things had to happen: the science team had to be able to operate remotely from their home institutions, and the 24/7 on-Mars-time approach had to be replaced with an approach involving weekday shifts on Earth time (Mishkin and Laubach, 2006). To facilitate this shift, each planning day had to plan between one and three sol's worth of rover activities, resulting in needing to squeeze even more effort into a planning shift. Fortunately, experience gained in commanding the rovers has resulted in adjusting the planning tools and their interactions to reduce effort, making the generation of up to three sol's worth of commands in a single day feasible. Even with these adjustments, moving off of Mars time has reduced how much the rovers do over a given week. The reduction takes the form of a constraint that a rover can only drive and move its arm once per operations shift. When a shift plans for multiple sols, the rover can only drive once in the multiple-sol period instead of every sol, but it can still use its instruments to collect science data on the other sols.

In the following section, this paper describes the issues underlying the strategic planning process, which sets the stage for the daily tactical operations process. The next section describes how the result of strategic planning combines with recently delivered data to determine the goals for the rover over the next one to three sols. The subsequent sections respectively describe activity planning, rover motion planning, and the integration/test of all sequences in preparation to the next uplink. Finally the paper summarizes and concludes.

To illustrate the interacting planning processes, the Integration Definition for Function Modeling (IDEF0) formalism (FIPS-PUB-183, 1993) is used. This is a simple graphical formalism where arrows from different sides of a box represent different things. Figure 2 essentially describes the graphical syntax, where inputs involve information feeding automatically into a process, controls are things to keep in mind while performing a process and a mechanism is some editing tool that supports the process.

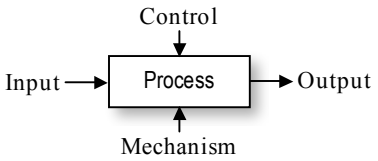

Figure 2. Graphical syntax for IDEF0 formalism.

Strategic Planning

The objective of strategic planning is to set the stage for daily operations. Strategic planning is a continual process that is decoupled from daily tactical operations. As such, it does not have any hard deadline other than being able to inform the tactical planning process of when each operations shift occurs and regarding the goals for the next few operations shifts. While the occurrence of an operations shift is determined months in advance, the actual goals have to adapt to rover uncertainty, and this adaptation takes the form of pushing goals into the future as problems occur or unexpected measurements result in motivating additional scientifically interesting observations. If a goal becomes infeasible due to conditions on Mars, it may be removed.

Communication Windows

Since orbital mechanics dictate communications windows, the existence of possible communications windows is know months in advance. For instance, we can only transmit commands to a rover after the Earth rises over the local Martian horizon, and there is a desire to send commands to a rover early in its morning to facilitate getting results from the next orbiter over-flight. Thus the commonly desired uplink time is after 9:30 (Mars local time) when the Earth is above the horizon and the solar panels are providing power.

While communications with Earth for uplinks depends on the relative positions of Earth and Mars, downlink times depend on the relative position between a rover and an orbiter. It turns out that both the Odyssey orbiter (ODY) and the Mars Reconnaissance Obiter (MRO) are in polar orbits about 450 km above the Martian surface. Thus they each provide two over-flights a sol, and each over-flight can provide a link of varying duration depending on how high the orbiter rises above a rover's local horizon. This results in the rovers having a communications window schedule that looks like the following in Martian local time.

02:00±0:30 – MRO AM over-flight
03:00±0:30 – ODY AM over-flight
11:30±2:00 – Uplink time
14:00±0:30 – MRO PM over-flight
15:00±0:30 – ODY PM over-flight

Given this daily schedule, communications window planning is mainly a matter of choosing the best windows from what is possible. Normally, the chosen windows include daily uplinks and PM relays with Odyssey. This maximizes time for performing activities with uncertain outcomes and responding to these outcomes in the next uplink. While AM passes are available, they happen during the Martian night, resulting in leaning heavily on the batteries. Thus they are rarely used, but a need to clear onboard memory occasionally motivates them.

Finally, the actual inclusion of each communications window involves a negotiation with the MRO, Odyssey, and the Deep Space Network. Each can either shorten or remove a window as needed by external constraints, and an adjustment can occur at any time.

Planning Shifts & Staff Assignment

While the timing of communications in a rover's local Martian time does not vary all that much, it does shift 40

minutes later every day on Earth. This results in there being times when a planning shift on Earth occurs during the same time that a rover is able to perform a plan. Thus the relationships between operations shifts and rover sols typically look like that in Figure 3, where the critical results of a drive or IDD motion must be downlinked before the next operations shift starts. In this case Opportunity has an operations shift every day of the week, but there are shifts for Spirit on Monday, Wednesday, and Friday only. Since the rovers are on opposite sides of Mars, the general pattern is that one rover will have daily operations to plan single sols while the other requires planning for two sols every other day. In both cases, Friday shifts involve planning for three sols to cover the weekend with rare two-sol exceptions.

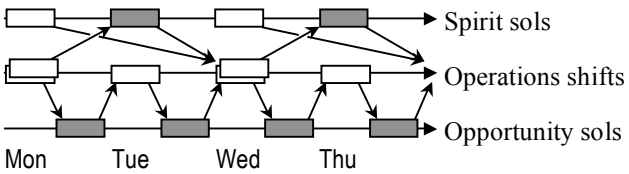

Figure 3. The critical data communications patterns between operations shifts and sols on Mars with allowed drive/IDD sols in gray.

Given this pattern of communications, when each shift occurs for each rover is known months in advance, and the problem becomes a matter of staffing each shift. This is a general job-shop scheduling problem where each shift needs six uplink engineering positions filled and the uplink sequencing team has a pool of people who have expertise in one or more positions. Finally, each team member has restrictions on how many shifts they can fill per week and specific dates where they are unavailable.

While there are tools for solving job-shop problems like this, this one is solved manually on an Excel spreadsheet every two months in order to get a feel for which positions need more trainees. There have been a number of staff changes over the past five years on MER. As people leave, others are brought in and trained to fill different positions. Manually solving the job-shop problem results in getting an understanding of which positions need more trainees.

Long-Term Science Planning

All through the mission a long-term science plan is maintained for each rover. This plan primarily involves managing desired science campaigns from day to day. For instance, a past goal for Opportunity was to drive to Victoria Crater and descend into the crater and analyze the bedrock at key points. While pursuing that goal, Opportunity also made numerous measurements for analyzing the Martian atmosphere. In general, long-term science planning is performed in team meetings and results in recommending a number of observations and drive activities over the next few operations shifts, with an eye to maintain a tradeoff between making additional observations to analyze some discovered feature and pushing on to reach some long-term science goal.

As hinted at by the IDEF0 in Figure 4, long-term planning is quite free form with data/pictures previously downlinked from the rover and the primary editing tool is PowerPoint with presentations to the tactical team at the start of every shift. The constraints on the process are the goals and onboard storage limitations. Finally the resultant outputs to a day's operations is the "sol path" denoting the recommended types of activities for the day, and a limit on how much data the desired activities can collect.

Determining the Day's Activities

The first step in a tactical operations shift is to determine the day's activities. This involves taking the sol path's extremely high level directive like "perform remote science" and actually figure out what engineering activities and science observations to execute. This determination involves taking as many observations as possible in the face of resource limitations revolving around energy, thermal, downlink, and data storage constraints.

As illustrated in Figure 4, determining a rover's activities and observations for the day involves an interaction between several processes. First numerous inputs are fed into skeleton planning, which subsequently constrains long term planning with an expected downlink volume. Activity window constraints from skeleton planning are then combined with a recommended data limit to subsequently constrain science planning, which determines the activities to perform.

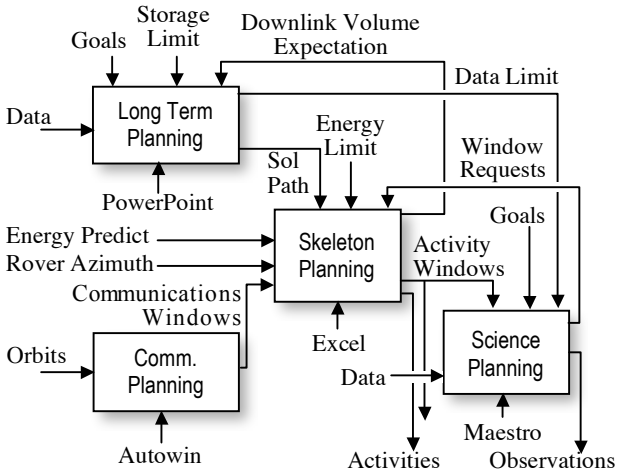

Figure 4. IDEF0 representation of how different abstract planning processes combine to determine the desired science activities.

As in most science planning problems, the issue is to maximize science collection in the face of uncertainty and resource limits. In the case of the MER mission, the tightest resource constraints are available power as well as available onboard storage. Each of these resources has a

certain amount of uncertainty. Available storage depends on how much data is transmitted in a given orbiter pass. The geometry of an orbiter's trace across the Martian sky interacts with the direction that the rover is pointed to effect how much data gets transmitted in a given pass. So how much data gets downlinked is unknown until it actually reaches the ground. Available energy also has a certain level of uncertainty due to dust in the atmosphere blocking sunlight, which is weather dependent.

Skeleton Planning

Skeleton planning is extremely high level and is implemented as a set of macros in Excel that builds a template spreadsheet to support an editing session. The first step is to copy in communications window information, expected rover azimuth, and solar energy predictions. The sol path is then manually used to determine which lines of the spreadsheet actually apply to the current sol's plan, resulting in determining activity windows and engineering activities. Next the equations underlying the start, end, and duration cells are checked. Finally, the durations of lines like "Remote Sensing" and "AM Science" are adjusted to come up with an acceptable net energy prediction for the sol (see Figure 5).

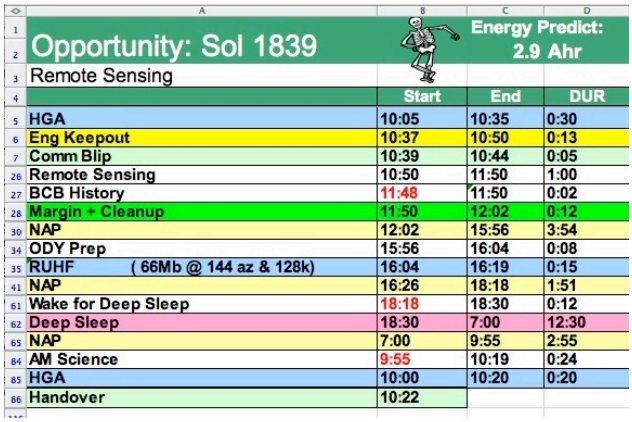

	A	B	C	D
1	Opportunity: Sol 1839		Energy Predict:	
2			2.9 Ahr	
3	Remote Sensing			
4		Start	End	DUR
5	HGA	10:05	10:35	0:30
6	Eng Keepout	10:37	10:50	0:13
7	Comm Blip	10:39	10:44	0:05
26	Remote Sensing	10:50	11:50	1:00
27	BCB History	11:48	11:50	0:02
28	Margin + Cleanup	11:50	12:02	0:12
30	NAP	12:02	15:56	3:54
34	ODY Prep	15:56	16:04	0:08
35	RUHF (66Mb @ 144 az & 128k)	16:04	16:19	0:15
41	NAP	16:26	18:18	1:51
61	Wake for Deep Sleep	18:18	18:30	0:12
62	Deep Sleep	18:30	7:00	12:30
65	NAP	7:00	9:55	2:55
84	AM Science	9:55	10:19	0:24
85	HGA	10:00	10:20	0:20
86	Handover	10:22		

Figure 5. Excel spreadsheet with a skeleton plan computing expected energy increase/decrease.

Initial skeleton plans are prepared the day prior to a tactical planning shift, and each shift has one to three skeleton plans depending on how many sols are being planned for. At the start of the day, these plans are presented and adjusted to satisfy requests to add and alter activity windows from the science team during science planning. Such alterations are rapidly made and tested by virtue of the set of equations in the original template of each excel spreadsheet.

Skeleton planning is an example of one of the larger streamlining improvements to the daily operations process. Initially the activity planning for a single sol took up to 4 hours. Generating a skeleton plan for a sol often takes less than 15 minutes, and these plans serve two purposes: they give the participating scientists a sense of how many observations they can request and they accelerate the activity planning process to the point where it often takes less than 45 minutes to plan for up to three sols. This improvement was only possible after months of operations experience resulted in determining the most commonly occurring high-level plan features and their timing relative to when communications windows occur.

Science Planning

Science planning involves assessing data from previous sols' observations to determine a rover's local environment and create new observations to forward mission goals in the context of new discoveries. As such, new observations are subject to a number of constraints. First of all, they must supply the data critical to determining the local environment for the next planning shift, and that data volume must be small enough to be downlinked prior to the next shift. Also, the total amount of critical and non-critical data collected must fit in the onboard storage, and the amount of time required to perform the observation and movement activities must fit in to the activity windows.

Science planning uses the Maestro science activity planner (Norris et al., 2005) shown in Figure 6, which has facilities for assessing local environments and modeling how long particular observations take, energy they consume, and how much data they generate. With all this information, participating scientists for each instrument take the activity windows and data limit and suggest lists of observations to fit into the activity windows and generate the recommended data limit's worth of data.

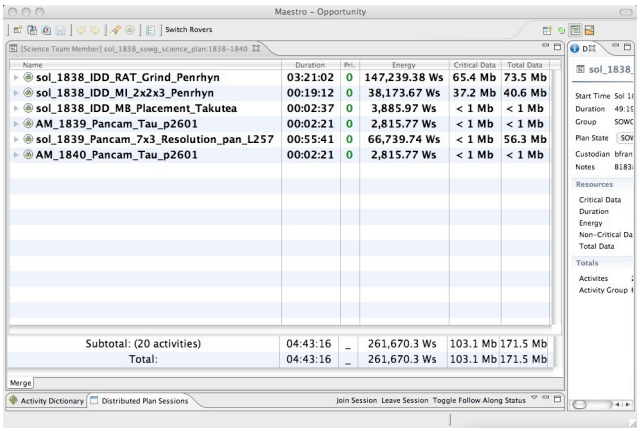

Figure 6. The Maestro science activity planner is an interface that the science team uses to inspect Martian imagery, specify science targets, define activities, and estimate time and data resource usages.

While generating a science plan, the need for some particular observation that does not fit into the windows provided by the skeleton plan might arise. In this event, a request is made to either add or modify a window in the skeleton plan. This requested change is manually applied to the skeleton and accepted if power permits. Thus skeleton planning and science planning interact to

determine the day's activities in the face of hard and soft time/resource constraints.

Planning/Scheduling the Day's Activities

Once the day's activities and activity windows have been determined, the next step is to plan/schedule the activities. As shown in the IDEF0 of Figure 7, planning/scheduling starts by constraining the science activities with respect to each other and the activity windows. This is done with the constraint editor (CE) (NASA/Ames HCI Group, 2004). While CE is quite powerful, most of the time the science activities are simply constrained to execute one after the other in their desired activity windows, which is nearly completely automated with a little manual checking/repair to detect and resolve the occasional rare error that slips in during skeleton and science plan editing.

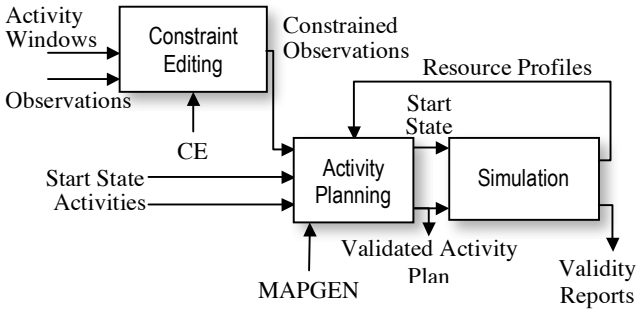

Figure 7. IDEF0 representation of how different activities interact during master planning/scheduling.

After the constraints are determined, all activities are loaded into MAPGEN, a mixed-initiative planner (Bresina et al., 2005). The activity planning and scheduling performed in MAPGEN involves a number of processes: expanding science activities to model their components, computing when the rover CPU must be on, adding heating activities, inspecting engineering activities to fix parameter errors, and altering activities to satisfy resource constraints. As shown in the IDEF0, MAPGEN feeds an activity schedule to a simulation, which returns resource usage profiles like the one illustrated at the bottom of Figure 8. The objective is to alter the plan as little as possible in order to get resource profiles within acceptable limits.

Plan manipulation primarily takes the form of deleting/shortening science activities to raise resource profiles to acceptable levels, and sliding other science activities earlier to fill gaps while satisfying timing constraints. This is the main reason for adding timing constraints among science activities in the CE, and it makes performing small changes quite painless, but there is a price. The addition of constraints makes larger changes more painful. A MAPGEN plan can involve more than a single sol's worth of activities. On occasion there is a desire to move a science activity from one sol to another, but such large changes violate established constraints. It is possible to make such a change, but difficult. Due to this difficulty and time constraints in operations, such changes are rarely requested. Most of the time, skeleton planning is precise enough to make MAPGEN planning a simple plan inspection process – resulting in constraint editing and activity planning requiring approximately half an hour to generate an acceptable activity schedule.

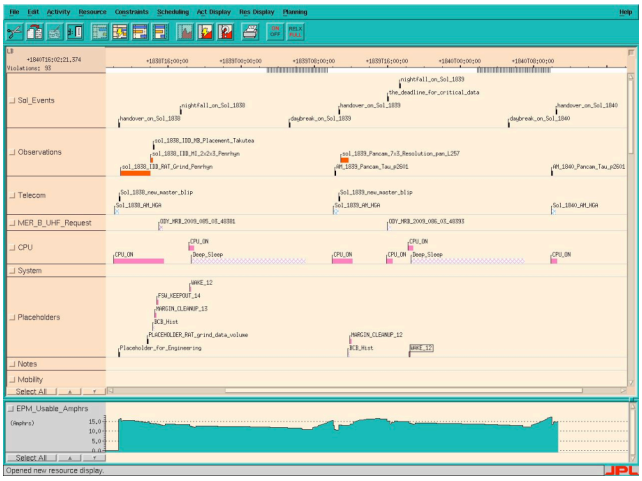

Figure 8. The MAPGEN planner with the activities appearing above a resource profile showing how available power evolves over time.

Once an activity schedule is generated, it is loaded into an Excel spreadsheet editor for annotation. While MAPGEN works with observations and activities, a rover executes sequences of commands. This mapping from an activity to a sequence requires some manual annotation due to a need to insert commands to resolve issues not modeled in MAPGEN. Figure 9 illustrates the processes underlying the translation of a MAPGEN schedule to master and submaster sequences, which call the sequences that implement science and engineering activities. As shown, a translator takes the annotated Excel schedule and builds the master and submaster sequences. These sequences are subsequently inspected in the Rover Sequence Editor (RoSE), which appears in Figure 10. In general, RoSE is a textual command, parameter, and comment editor used to edit all rover sequences, including the creation of the sequences that implement planned activities.

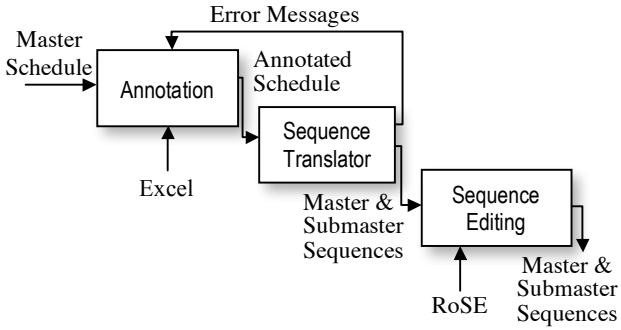

Figure 9. IDEF0 of how processes combine to translate a schedule into valid master and submaster sequences.

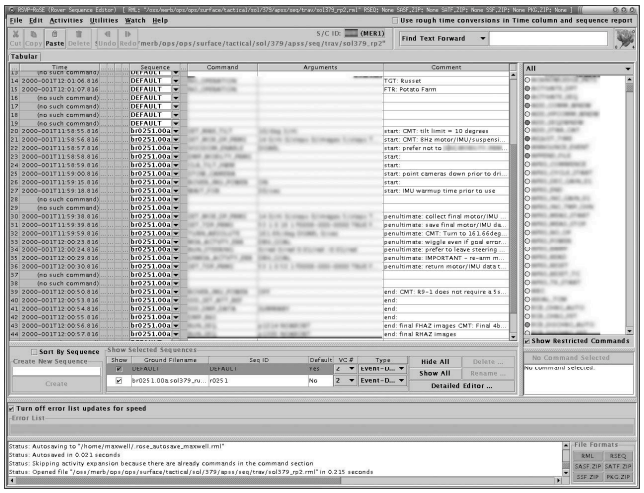

Figure 10. Rover Sequencing Editor (RoSE) with parts blurred due to legal restrictions.

Mob/IDD Planning

In parallel with generating the (sub)master sequences, sequences for implementing rover motion are crafted. These motions involve driving from one location to another (Mobility) and moving the Instrument Deployment Device (IDD), which is a robotic arm. Where the (sub)master sequences are clock and event driven to be highly deterministic, Mob/IDD sequences involve highly uncertain activities and have to deal with this uncertainty using conditional sequences and high-level sensor-based commands for traverse. For instance, a sequence to drive to a certain location while avoiding known obstacles may seem straightforward, but there are a number of hidden problems. There is no way to know what lies below the surface dust being driven upon. So a rover can slide in unexpected directions as it drives, and can even get stuck by digging a wheel into a hole. While some of these contingencies are handled by artificially intelligent flight software (Maimone et al., 2006), dangerous uncertainties still creep in while executing these smart commands. For this reason rover planners need to break drives up into small steps that alternate with progress/safety checks, which in turn leads to sequences with conditional execution. Ultimately these sequences tend to be quite complex, and numerous visualization tools are used to assist the Mob/IDD planning effort.

As shown in Figure 11, the two main processes that contribute to building Mob/IDD sequences involve textually editing sequences in RoSE and graphically visualizing/changing sequences in HyperDrive (Maxwell et al., 2005; Yen et al. 2005; Write et al., 2005). Sequence generation is primarily a manual activity that draws from past sequences/experience, but each sequence must be tailored as local hazards change when a rover moves away from one set of hazards and toward another. Thus HyperDrive provides an immersive visualization for viewing and specifying geometric components of a Mob/IDD sequence, and RoSE provides a textual interface for adding conditional execution branches to respond to uncertainty.

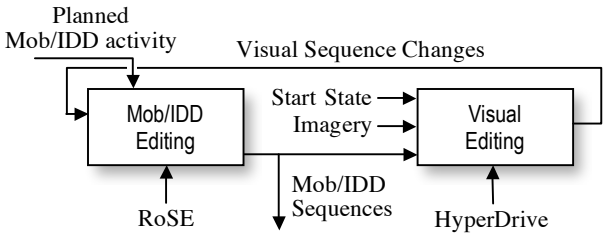

Figure 11. IDEF0 of how textual and visual editing processes combine to build a Mob/IDD sequence.

To provide a visualization of a rover in its local environment, HyperDrive needs a number of stereo images from the rover. These are downlinked as critical data to derive a 3D terrain map of a rover's location relative to obstacles on Mars. Figure 12 gives an example visualization of Opportunity with respect to its landing site and the bedrock first imaged upon landing. In addition to Figure 12's image, the tool lets a Mob/IDD planner manipulate the visualization to look at the rover in its local environment from any angle, including a simulated image from one of its cameras. This immersive visualization of a rover's current state as well as how the rover would move if it executed commands has proven to be a powerful tool for crafting and debugging Mob/IDD sequences.

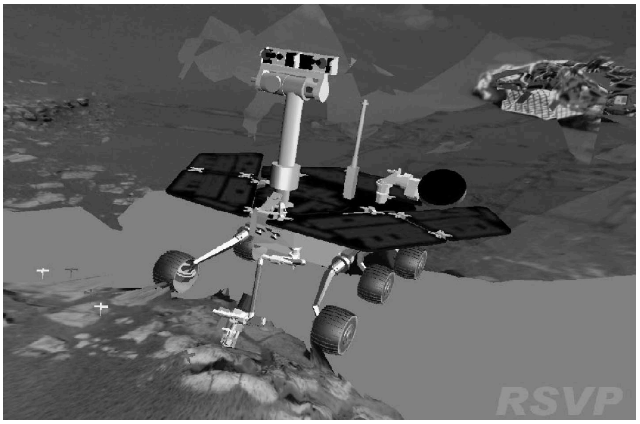

Figure 12. Sample HyperDrive visualization. (Image courtesy of Frank Hartman, JPL)

Plan Integration & Test

While the master, submaster, and Mob/IDD sequences are developed at JPL, many of the sequences that implement instrument observations are often generated remotely at the home institutions that built the commanded science instruments. As mentioned previously, each rover has seven science instruments. Each instrument can be

commanded with a number of different sequences to implement a variety of observations. While the observations were determined in science planning, their implementing sequences still needed to be crafted and delivered. Even this process has been streamlined over time through the use of collected experience. The MER mission maintains an archive of all observation sequences that have ever executed on Mars, and getting a sequence to implement an observation is often a matter of selecting it from the archive. It turns out that many observations have multiple implementing sequences, and determining which sequence to use depends on a number of circumstances including target angles relative to the rover, time of day, lighting conditions, relative importance of the collected data, and others. The number of circumstances is large, resulting in still needing human assistance to determine the sequences to include, but the effort in implementing observations is reduced to either just a selection or a selection and a few adapting edits in RoSE.

Once all sequences have been determined, the next step is to integrate them together and test them for global consistency. As implied by the IDEF0 in Figure 13, the integration involves running automated scripts without any need for an editing mechanism. The resultant integrated sequence is then loaded into RoSE, and a batch simulation checks it for timing errors and flight rule violations. In the rare event that a violation is found, either the violation is waved or a final edit followed by another simulation is performed to repair the problem. After all violations have been accounted for, the integrated sequence is ready to be sent to the rover.

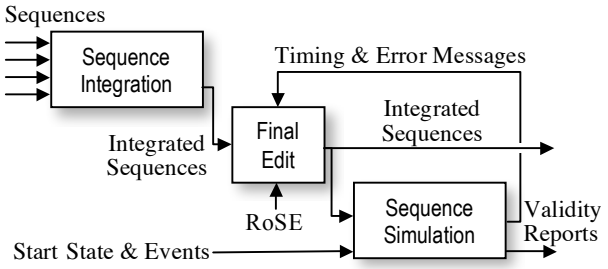

Figure 13. IDEF0 representation of the processes that contribute to generating the final integrated sequence.

Conclusion

This paper described the planning/scheduling processes that go into building an integrated sequence for a MER rover in a single day's operations. The IDEF0 characterizations of processes illustrated in previous figures link up as shown in the IDEF0 representation of Figure 14.

Globally, the inputs to the day's operations are initial conditions & events, which include all the information needed to determine a rover's status, its local environment, and exogenous events that interact with a rover's actions. The goals and resource constraints combine to determine what the rover can and should do. Finally, the outputs of the process are a set of integrated sequences to be transmitted to the rover as well as a number of validity reports that attest to the sequences' correctness.

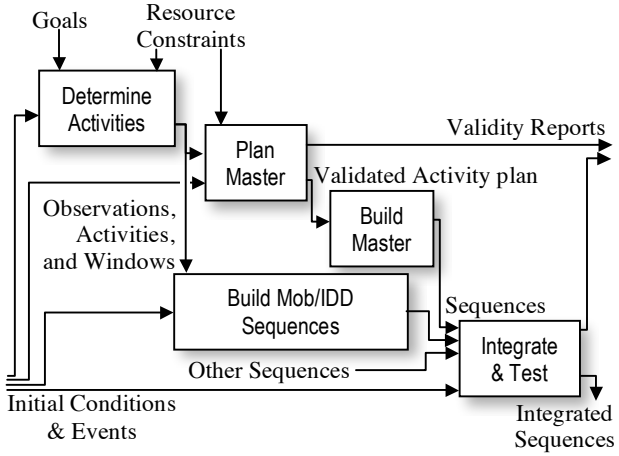

Figure 14. IDEF0 representation of the daily planning process for MER project.

The operations day starts with a team-wide Science Operations Working Group (SOWG) meeting to determine the rover activities, with an objective to craft integrated sequences that command a rover for one to three sols. After this meeting the team separates to craft sequences in parallel. The team reconvenes for another meeting to inspect/approve the planned master schedule, its validity reports, and the visual Mob/IDD motions. After this meeting the sequences are completed and another meeting convenes to inspect/approve the sequences prior to integration. Finally, after integration & test, the final integrated sequences are inspected/approved for uplink.

While the tools take much of the drudgery out of a day's operations, the uplink planning/sequencing activity is still intensely manual with low-level inspections of each newly crafted sequence. On the one hand this manual process provides multiple opportunities for human error, but on the other hand it also provides multiple mechanisms for recovering from error as well as adapting to situations like surviving a conjunction, a winter, or a dust storm.

While there were a number of lessons learned over the past 5 years, the three below are most relevant to the planning and scheduling community. These lessons are not necessarily new, but they have been ignored more than once.

Tool interfaces should be as simple as possible, commensurate with the task at hand

This lesson drove the addition of an Excel-based skeleton planner. When performing initial science planning, an operator is only interested in awake times and power utilization. Working with the full-up activity planner imposed unnecessary overhead—and therefore time—for

the relatively high level and low precision initial planning performed at the start of the tactical operations process.

Rover capabilities and operational strategies both evolve over time

From the occurrence of a solar-array-cleaning gust of wind to the freezing of a robot arm joint, rovers change and age over time, and how they are operated must change in response; team experience and budget constraints also encourage continuing efforts to optimize the operations process. Tools must be sufficiently flexible to accommodate these changes whenever possible. For example, manual annotation of the MAPGEN-produced master activity plan enabled the use of scripts to automatically generate the master and submaster command sequences that implement that plan. These sequences include housekeeping commands that were not tracked by the planner. As the rovers aged and associated operational mitigations were developed, new commands and housekeeping sequence calls were introduced into the master and submaster sequences. The use of new annotations and updated scripts to support them has proven to be a flexible means of ensuring consistent sequence products.

Sometimes flight rules can be waived

In the planning and scheduling community, flight rules are often modeled as hard constraints, and violated constraints are flaws that a tool repairs. During actual flight operations, flight rules may at times be formally waived, under limited conditions, e.g., if situational awareness and analysis of the current circumstances determine that the rule is not applicable, or if the spacecraft is already at risk due to an anomaly, and an action that would violate the flight rule has been determined to be a means to reduce that risk. The usefulness of tools may be significantly reduced if they do not allow—with appropriate safeguards—for such eventualities.

Acknowledgements

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (NASA). Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

References

Bresina, J., Jónsson, A., Morris, P., and Rajan, K. (2005) “Activity Planning for the Mars Exploration Rovers”, In Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling, Monterey, 2005, pp. 40-49.

FIPS-PUB-183. (1993) “Integration Definition for Function Modeling (IDEF0).” National Institute of Standards and Technology.

NASA/Ames HCI Group (2004) “Constraint Editor”, http://hci.arc.nasa.gov/pages/2004/10/constraint_edit.html.

Maimone, M., Johnson, A., Cheng, Y., Willson, R., and Matthies, L. (2006) “Autonomous Navigation Results for the Mars Exploration Rover (MER) Mission.” Springer Tracts in Advanced Robotics, Volume 21. 2006.

Maxwell, S., Cooper, B., Hartman, F., Wright, J., Yen, J., and Leger, C. (2005) “The Best of Both Worlds: Integrating Textual And Visual Command Interfaces for Mars Rover Operations.” In Proceedings of IEEE/SMC 2005.

Mishkin, A., Limonadi, D., Laubach, S., and Bass, D. (2006) “Working the Martian Night Shift: The MER Surface Operations Process.” In IEEE Robotics Automation Magazine Special Issue on Mars Exploration Rovers.

Mishkin, A. and Laubach, S. (2006) “From Prime to Extended Mission: Evolution of the Mars Exploration Rover Tactical Uplink Process.” *Proceedings of AIAA SpaceOps Conference.*

Norris, J., Powell, M., Vona, M., Backes, P., and Wick, J. (2005) “Mars Exploration Rover Operations with the Science Activity Planner.” *Proceedings of 2005 IEEE International Conference on Robotics and Automation*, Barcelona, Spain. April 2005.

Wright, J., Trebi-Ollennu, A., Hartman, F., Cooper, B., Maxwell, S., Yen, J., and Morrison, J. (2005) “Terrain Modelling for In-Situ Activity Planning and Rehearsal for the Mars Exploration Rovers.” In Proceedings of IEEE/SMC 2005.

Yen, J., Cooper, B., Hartman, F., Maxwell, S., Wright, J., and Leger, C. (2005) “Physical-Based Simulation for Mars Exploration Rover Tactical Sequencing.” In Proceedings of IEEE/SMC 2005.