Driving on the Surface of Mars with the Rover Sequencing and Visualization Program

J. Wright, F. Hartman, B. Cooper, S. Maxwell, J. Yen, J. Morrison
Jet Propulsion Laboratory, 4800 Oak Grove Dr., Pasadena, CA, USA
John.R.Wright@jpl.nasa.gov

Abstract: Operating a rover on Mars is not possible using teleoperations due to the distance involved and the bandwidth limitations. Instead, rovers are autonomous and go about their operations in an unsupervised mode. To operate these rovers requires sophisticated tools to make operators knowledgeable of the terrain, hazards, features of interest, and rover state and limitations, and to support building command sequences and rehearsing expected operations. The primary requirement is to rapidly convey the optimal level of situational awareness to the operators to support daily cycles of operations and commanding. This paper discusses how the Rover Sequencing and Visualization Program and a small set of associated tools support this requirement.

Keywords: Mars, rovers, autonomy, commanding, ground operations, telepresence.

1. Introduction

The Jet Propulsion Laboratory has been operating twin rovers on Mars since January of 2004. The two rovers, named Spirit and Opportunity and part of the Mars Exploration Rovers (MER) mission, are exploring Gusev Crater [LEGER05] and Meridiani Planum [BIESIADECKI05] respectively. Because the rovers are solar-powered, their operational cycles are tied to the Martian diurnal cycle. Consequently, Earth-bound operators are also tied to the Martian day, or sol, which is about 40 minutes longer than an Earth day. The similarity in the length of the day on the two planets, as well as the desire to maximize science return while also maintaining the safety of the rovers, leads to several operational constraints. These constraints combine to essentially mandate a daily mission planning cycle in which downlinked data is processed, the rover state is analyzed, the terrain is studied for hazards, traversability, and features of interest, and new commands are uplinked every sol.

The Rover Sequencing and Visualization Program (RSVP) is a set of software tools that are designed to work in this daily planning environment to support the rapid analysis of rover state information and creation of command sequences. Previous planetary missions could plan for encounters known months or years in advance. The Mars Pathfinder mission pioneered the use of the daily planning cycle [COOPER98] and this paradigm was extended for MER. The RSVP tools include terrain model visualization and interaction, numeric data plotting and analysis, image display and interrogation, command sequence visualization, sequence rehearsal, kinematic modelling of rover and terrain interactions, and time-based modelling of spacecraft and planetary bodies for analysis of communication issues, incident solar energy, and shadowing.

Essentially, RSVP must provide the operators with information on the local terrain and rover state in a mode that supports rapid assimilation and understanding. Then it must support rapid assembling and testing of command sequences for correct execution. It must also support the creation of documentation and archival products.

2. Data Analysis

RSVP ingests several types of data to support its mission. These data include current and historical state information, imagery with associated state information, three-dimensional terrain models, and science activity plans.

The state information is used to analyze and review the current state of the rover, identify any anomalous issues, review previously commanded activities, and verify that the rover is ready to accept and perform a new set of commanded activities. Imagery is used to visually examine the surrounding terrain, identify hazards and areas of interest, specify targetting information for additional observations, and essentially present the operator with the least processed view of the current environment. The terrain models are used similarly but with full three-dimensional representations of the terrain and unlimited viewing from any position and orientation. The science activity plans are produced by the science team and are used as guides for planning approaches to targets of interest and for defining desired activities for the instruments on the Instrument Deployment Device (IDD), the robotic arm.

2.1 State Analysis

The state of the rover is represented by numeric values collected by numerous sensors located at various points on the rover, as well as the results of internal calculations based on the sensor data. Examples of data channels (the data associated with a single sensor) include the rocker and bogey joint angles of the suspension, the wheel steering angles, the wheel rotation counters, the joint angle sensors on the arm, temperature sensors, and many others. Examples of internally computed channels include the XYZ location of the rover via dead reckoning or visual odometry.

Analysis of the state information takes several forms. The first is a simple display of all the numeric values, organized into several pages of displays, each focussing on a particular subsystem such as mobility or IDD. There are also summary pages that present the most pertinent information for multiple subsystems. This type of display has only limited usefulness because it represents only a single snapshot in time (typically the latest or current rover state) and is so full of detailed information that it is difficult to locate pertinent data without a great deal of experience. Some functions do aid the operator by color-coding particular pieces of information, generally red to indicate an error condition or blue to indicate a value that has changed from the last update. However, this type of display is mostly useful in identifying a particular subsystem with an anomalous condition to aid in focussing on a particular problem. For example, for several weeks during the mission, the rover Spirit was having trouble with its right front wheel. The problem manifested itself in excessive current draw by the drive motor on that wheel and this was shown on the numeric displays immediately. However, it took examination of time histories of current draw over several weeks to track the progression of the problem, determine probable life expectancy of the motor, and suggest mitigation activities. One such activity was to drive the motor back and forth while warming it with its internal heater in an attempt to redistribute the lubricant. This had minimal effect so usage of that drive wheel was curtailed in order to extend its life. The rover was driven backwards while alternately dragging and driving the wheel, with a 10% duty cycle, to maximize the range available. This method was used successfully for several weeks until it was noticed that the current draw had dropped back to normal and regular usage of the wheel was authorized.

From the previous example, it can be seen that the use of time histories of state information is critical to analyzing many issues. Many types of plotting tools are available and RSVP includes a set for analyzing a variety of mobility issues. The channels that can be displayed in RSVP pertain almost entirely to mobility and arm operations and include XYZ position, orientation, overall tilt to investigate tilt limit issues, northerly tilt to examine incident solar energy issues, etc. RSVP makes no distinction between state information collected onboard the rover during operations and state information collected during simulation and rehearsal of planned operations. Thus, the data can be displayed side-by-side for rapid comparison of planned and actual activities of the rover.

State information from actual rover activities is typically used to analyze the behavior of the rover during traverses, especially when using its internal hazard avoidance system. The rover may arrive at its intended destination through a circuitous route that expends additional energy and puts additional wear on components when a simpler route might have been available. Understanding the rover's hazard avoidance software's interactions with various terrain types aids the operators in planning optimal traverses. A particular example was a planned traverse of about 40 meters across a rocky terrain. As the rover neared a cluster of rocks, it identified the rocks as hazardous and attempted to find a path to the left of the cluster. It drove forward and backward, slowly turning left, without finding a clear route to the left until it had turned far enough for the rear of the rover to be pointing to the right of the cluster. The rear cameras imaged the terrain, identified a clear route, and the rover proceeded backwards around the rock cluster and the rest of the way to the goal, only turning to face the destination at the last moment. Examination of final state would show that the rover reached the destination. Examination of the intermediate route would show the interesting behavior and suggest possible terrain types to avoid and preferential routing through problematic areas.

Arm activities of the rover rarely need to be analyzed this way because the flight software for controlling arm operations is integrated into RSVP, as discussed in more detail later. Thus, arm activities that successfully simulate in RSVP almost always work as intended in operations. The exceptions are due to inexact knowledge of the terrain, which leads to insufficient knowledge of actual arm-terrain interactions. To overcome this, a range of contact positions of instruments and terrain are simulated and commanding adjusted to avoid problematic areas.

2.2 Image Browsing

The MER rovers each carry a total of nine cameras or imagers, as shown in Figure 1. Each camera has a CCD array of 1k x 1k pixels with 12 bits of resolution per pixel. They are arranged in four pairs of stereo imagers, with the ninth being the Microscopic Imager as one of the instruments on the IDD. Each type of imager has a primary use and the field of view and mounting take that into account.

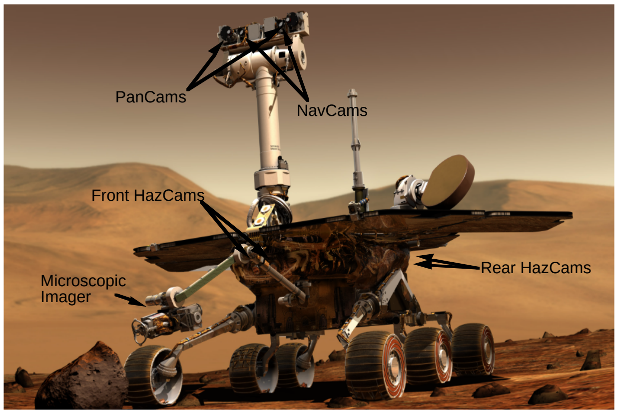

Figure 1 – Rover Cameras

The HazCams, or Hazard Avoidance Cameras, are mounted in fixed positions below the deck front and rear for examining the terrain within two to four meters of the rover's current location. They have a wide field of view and are tilted down such that the nadir is within the field of view. They are used both for analyzing the terrain for hazards and for imaging the terrain to be interacted with by the arm instruments.

The NavCams, or Navigation Cameras, are mounted on the mast above the deck and may be slewed and tilted to point in almost any direction. Each has a field of view of about 45° and are monochrome, like the HazCams. The NavCams are primarily used for imaging nearby terrain for three-dimensional modelling and traverse planning.

The PanCams, or Panorama Cameras, are also mounted on the mast above the deck. Each has a field of view of about 16° and contains a filter wheel for capturing multispectral imagery. The PanCams are primarily used for science observations in the multispectral domain but also may be used for modelling of terrain for long-range planning.

Each stereo pair captured by any of the imagers is processed on the ground to utilize the stereo information to compute the shape of the features being viewed. Generally, several

PanCam, NavCam, and HazCam image pairs are captured each sol during operations. This imagery may be reviewed and analyzed using the Image Browser tool of RSVP. The Image Browser reads in the specified imagery and displays it in context with other imagery, using camera model data stored in the image files to properly mosaic the images. Figure 2 shows an example of the display of the Image Browser with several NavCam images.

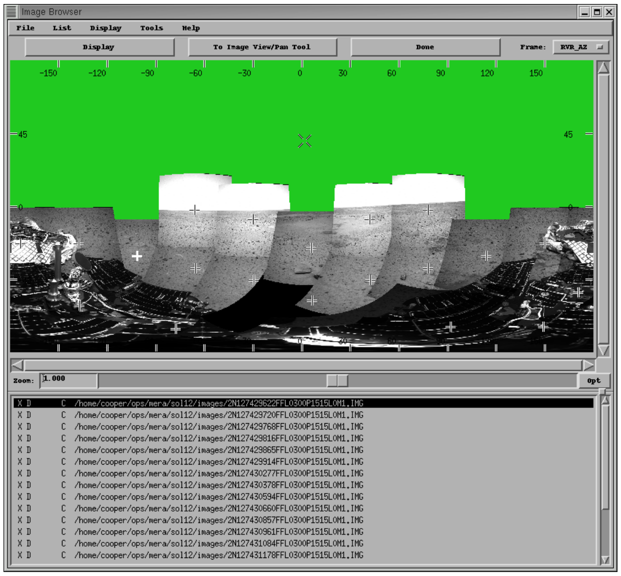

Figure 2 – Image Browser Display

When loading in images, the corresponding range and xyz images are loaded as well. This allows the Image Browser to be used to analyze range, slope, roughness, and other information for planning traverses. It also provides the capability of specifying a target and computing the surface normal. Many RSVP activities utilize targets specified by the Image Browser, or imported from the Science Activity Planner, for building command sequences.

In addition to the Image Browser, RSVP includes an additional image viewing tool, the Stereo View tool. This tool accepts image information from the Image Browser and displays the corresponding left and right stereo images, either individually or in stereo mode for use with CrystalEyes® glasses for an immersive effect. Use of the stereo glasses presents the operator with the least processed imagery with the highest fidelity for rapid understanding of the local terrain conditions. Small ridges that occlude hidden areas become quickly apparent in the stereo view and surface roughness and rock clusters are easy to understand. One twist is that the difference in the field of view of the PanCam, NavCam, and HazCam imagery presents challenges to the viewer to correlate what is seen with reality. Specifically, the narrow field of view of the PanCam gives the sense that faraway features are up close while the wide field of view of the HazCam does the opposite. Often the solution is to move the head closer to or farther from the screen such that the display window approximates the camera field of view relative to the eye.

2.3 Terrain Modelling

One drawback of the Image Browser and Stereo View tools is that their view is limited to that of the cameras on the rover. It is quite easy to tell if one rock is closer to the rover than another rock. It is much more difficult to see if the rover will fit between the rocks. Essentially, the operators need to be able to visualize the rover within the terrain and analyze interactions between the rover components and the terrain. Unfortunately, the rover cannot stand back and capture imagery of itself. To solve this problem, the range and XYZ information generated from the stereo imagery is used to build three-dimensional terrain models for visualization and simulation purposes as described in [WRIGHT04] and [WRIGHT05]. Then a virtual rover model can be used to examine terrain interactions. The terrain models contain two primary components, a multi-resolution, multiple level of detail triangle mesh, and a height map. The mesh is used for visualization while the height map is used for computing interactions between the wheels and the terrain. Figure 3 shows an example of a mesh built using NavCam imagery taken near Spirit's landing site in Gusev crater as visualized in RSVP.

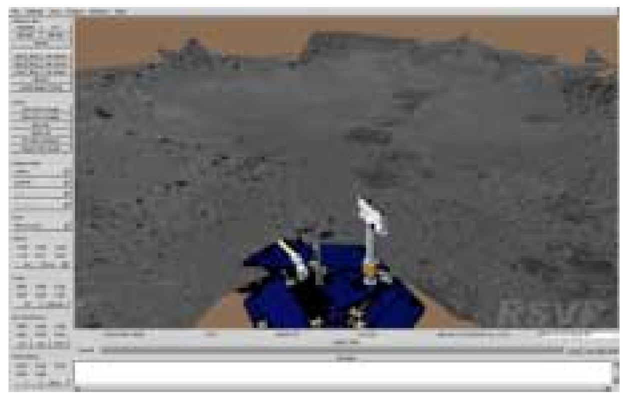

Figure 3 – Terrain Model Built from NavCam Imagery

2.4 Terrain Visualization

Visualization of the terrain models is performed by the HyperDrive tool, which has several modes. The primary use is the flying camera in which the operator moves a virtual camera around the terrain model using a tabletop or free motion mode. Bringing the camera up to look down on the terrain from above allows the operator to quickly identify areas hidden from the rover as these will be unpopulated with terrain data and will appear as holes in the terrain mesh. It also allows the operator to quickly get a big picture view of the surrounding terrain. The center of the tabletop can be the rover or a target to aid in visualizing the terrain around them.

Several additional visualization aids can also be displayed within HyperDrive. These include a horizontal plane to easily locate the origin and identify up, cross, and down-slope directions. JPL's SPICE kernels are integrated with HyperDrive allowing visualization of ephemeris data showing the direction to various bodies in space including the sun, Earth, Mars's moons Phobos and Deimos, and various orbiting spacecraft. In addition, HyperDrive can also display shadows cast by the sun at any desired time, allowing the operator to visualize lighting and shadowing of targets during planned observations. Icons for planned commands can be overlaid on the terrain models as well as rover tracks from simulations or actual telemetry.

Of great importance is the ability to visualize interactions between a virtual rover model and the terrain. HyperDrive contains a rover model that becomes a cursor in one mode. The cursor is dragged over the terrain and rotated to any orientation and the wheels, along with the rocker and bogey suspension, track the terrain model in realtime. The rover-terrain interactions are computed using the height map for easier sampling in the kinematic simulation. In addition, the rover cursor location may be used to specify targets or waypoints.

HyperDrive provides multiple views into the virtual world. One such view is the Stereo View tool described above and it allows the rover model, the command icons, the rover tracks, and other visualization aids to be displayed within the images seen in the tool. This is especially useful in stereo mode where the three-dimensional terrain is enhanced with three-dimensional representations of these aids, including the three-dimensional rover model. Figure 4 shows an example of the rover model overlaid on the imagery in the Stereo View tool.

Figure 4 – Stereo View Window (Mono Mode) with Rover

Another tool that visualizes the same environment is the Camera View window. This window displays the view from any of the cameras on the rover, including the Microscopic Imager on the IDD, from the current rover position and state. When planning observations with the cameras, it is quite useful to visualize what the camera is expected to see prior to the imaging. Often, this visualization is extremely accurate as seen in Figure 5. However, the accuracy is highly dependent on the quality of the terrain model in the vicinity of the feature or area being imaged.

3. Sequence Generation

Once the current state of the rover has been analyzed and the surrounding terrain reviewed for potential hazards and targets of interest, it is time to begin building the command sequences for the current sol's activities. The rovers are autonomous in that they conduct their activities while out of contact with ground controllers. However, the true level of autonomy varies with the types of commands being executed. The command with the highest level of autonomy is the go to waypoint command with hazard avoidance. This command performs terrain analysis, to identify hazards, and path selection to attempt to reach the commanded destination safely. This command is essentially nondeterministic because terrain knowledge on the ground is always inferior to knowledge onboard the rover.

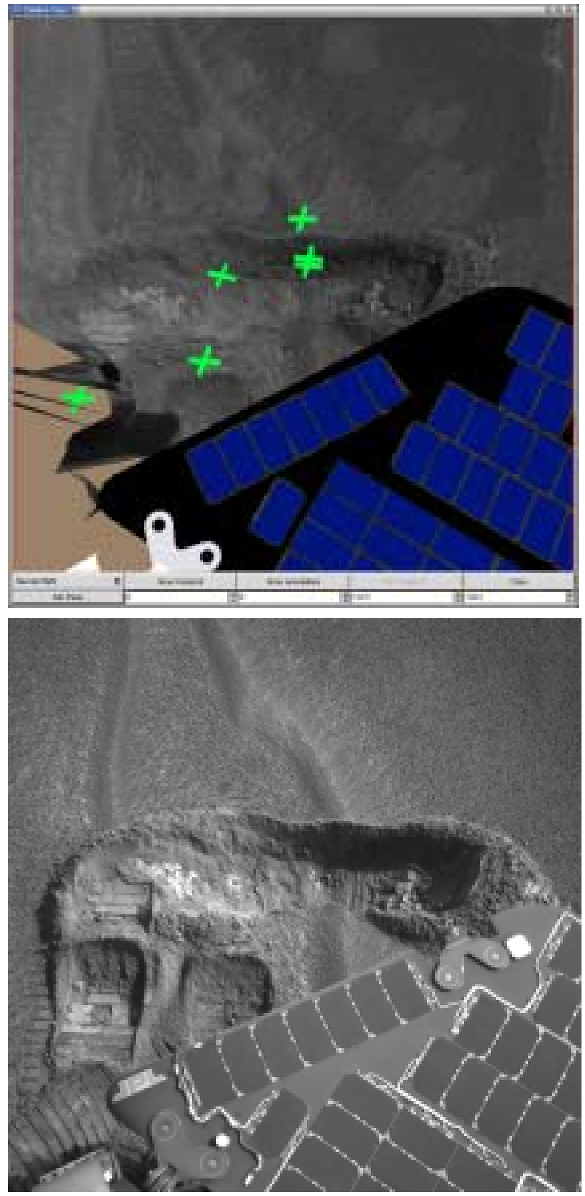

Figure 5 – Predicted (top) and Actual Camera Views

Commands are bundled into sequences that resemble subroutines in a program. Each sequence is typically dedicated to a single activity such as mobility, IDD, imaging a particular target, etc. The rover operators are responsible for the mobility and IDD sequences, dedicated Project Uplink Leads (PULs) for each instrument create the sequences for performing specific observations and controlling specific instruments, and the Sequence Integration Engineers (SIE) provide the master and submaster sequences for the sol. The master sequence for the sol is responsible for the overall timing of events, including shutdowns, restarts, activation of submasters, deactivation of other sequences, etc. The submaster sequences generally control the activities for a block of science observations, mobility, or IDD activities that perform consecutively.

Sequences can run sequentially or in parallel. In fact, the master sequence is always running and all other sequences run in parallel to the master. This allows the master to terminate other sequences if they take longer than expected to execute and other activities are required to begin.

3.1 Rover Sequence Editor

The Rover Sequence Editor (RoSE) is the backbone of the command sequence generation process. RoSE contains the master view of the sequence, maintains the unique identifiers for identifying selected commands in all the various tools, and distributes command and simulation updates throughout RSVP. Figure 6 illustrates the appearance of the RoSE interface. At its most fundamental, RoSE is a text editor for entering commands in sequence and for editing existing command sequences. However, RoSE provides many features that aid the operator in rapidly producing valid sequences. It is used by every operator, SIE, and PUL to produce the command sequences for delivery, documentation and archive products for review, and the final command bundles for uplink. Details outside the scope of this paper can be found at [MAXWELL04] and [MAXWELL05] but some of the highlights will be touched on here.

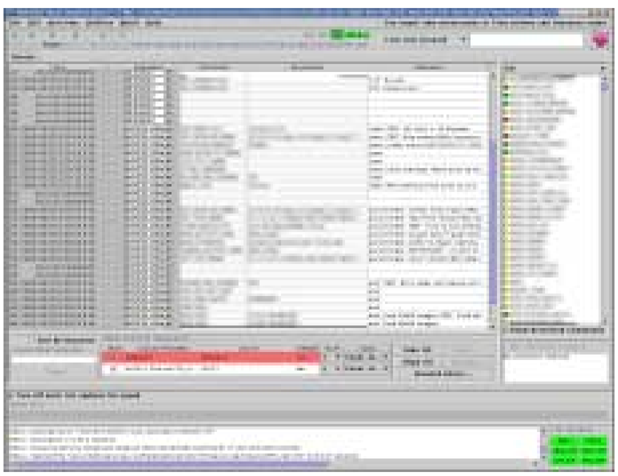

Figure 6 – RoSE User Interface

3.1.1 Macros

When building the actual command sequences, RSVP offers significant assistance for the rapid production of validated sequences. On Mars Pathfinder, the total size of the sequences sent to the rover were around 100 commands [COOPER98] per sol. On MER, the sequences can total over 1000 commands for a single sol and mobility and IDD sequences have exceeded 500 commands on occasion. Building and validating such large sequences requires support from the tools. Many sequences have similar blocks for performing functions that must be done each sol or for each activity. For example, prior to using the IDD, the temperature sensors on the arm joints must be tested and temperatures set to provide for proper current limit measurements. Since this must be done prior to all IDD activities and is done the same way each time, the commands for the process are stored in a macro and expanded when desired. Similarly, when ending the IDD activities, the recorded activity information, the state history for the activity, must be saved and then compressed for downlink in the same way every time. A macro encapsulates this behavior as well.

The macros provide a simple way to specify a few parameters and generate many lines of valid commands. In addition, the macros make it more difficult to miss a critical command in a large block of commands if the block was generated entirely from a single macro. The macros also provide templates for activities in which many commands are present and markers indicate a required additional command that must be generated separately. For example, IDD targets are specified in site frame while the IDD commands to access those targets are issued in rover frame. Rather than the user converting the coordinates, the HyperDrive tool is used to specify the desired site frame target, convert the coordinates, and generate the IDD commands. These commands are then inserted into the command block generated by the macro at the designated locations.

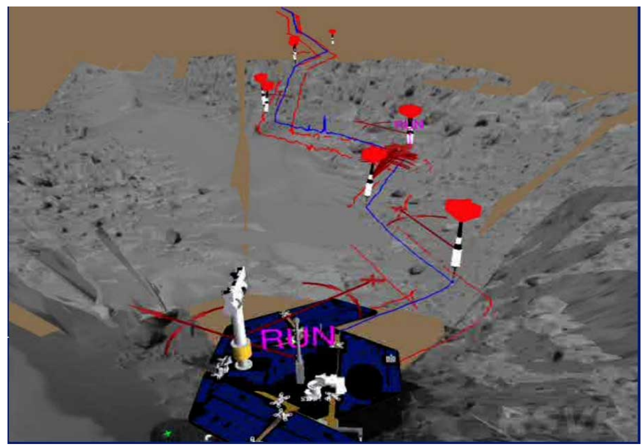

Figure 7 – HyperDrive Terrain and Icon Display

3.2 Sequence Completion

As mentioned above, macro expansion produces large portions of a sequence for most IDD and mobility activities. However, the specific motions that depend on target locations and directions, or on hazard avoidance, must be generated with knowledge of the terrain and the targetting information. These commands are generated using the HyperDrive tool. All the tools in the RSVP suite are tightly coupled with a shared view of the command sequence being developed. As previously mentioned, RoSE maintains the master view of the sequence and communicates this view to all the other tools when any change is made. Command insertion requests from HyperDrive are communicated to RoSE which performs the insertion and broadcasts the inserted commands to all the tools. While HyperDrive contains the entire command sequence under development, only the commands with a visual component are visualized. This would include waypoints, turns, and IDD movements but not heater commands, switch state testing, or other commands not coupled to the terrain. Figure 7 shows how HyperDrive uses a variety of icons to represent the commands and their relationship to the terrain.

3.3 Sequence Validation

Sequence validation is a multi-tiered process. Validation within RSVP ranges from syntax checking of commands and range-checking of arguments through simulation of the commands via software and on up to running the commands through the flight software itself or onboard an actual rover in the testbed. These steps verify that the sequence is syntactically correct and would execute on the rover. Flight rule checks verify that a variety of rules are not being broken. These might verify such things as that the dust cover on the Microscopic Imager is opened prior to taking an image. This provides additional feedback on the validity of the sequence. An additional level of check is the rehearsal and playback of the sequence for visual verification that the sequence performs correctly and interactions with the terrain are as expected. Finally, if some major component of the sequence has never been done before, it is often run on an actual rover in the JPL testbed to verify proper behavior prior to utilizing it on Mars.

Syntactic checking of the commands is performed continuously. RoSE ingests the command dictionary at startup time and builds the user interface for entering commands and checking parameter types and ranges automatically. Invalid commands or parameters are immediately highlighted to alert the user.

3.3.1 Sequence Simulation

Sequences are simulated almost continuously during development. As the user is adding commands to the sequence or editing or deleting existing commands, the sequence is sent to a tool called SEQGEN, developed at JPL and adapted to each mission and spacecraft. SEQGEN performs simulation of the sequence and returns the time at which each command is expected to execute, if ever, including reporting of how many times each subsequence was called and which commands executed during each subsequence execution. The durations of most commands are known a priori and these durations are used for the simulation. However, SEQGEN does not have information on the terrain or the rover state so it is unable to correctly estimate the duration of mobility and IDD motion commands. This part of the simulation process is performed by HyperDrive.

3.3.2 Sequence Rehearsal and Playback

HyperDrive performs simulation of sequence execution through links to actual flight software for IDD activities and through kinematic simulation tools for mobility activities [YEN05]. Once the simulation has completed, it can be played back to visually verify the proper behavior of the sequence. A VCR-like control widget is brought up and used to control the playback of the sequence as an animation in the HyperDrive window. As the rover model performs the commanded activities, the operator can view the animated model from any angle to verify that the behavior is as expected. The operator can stop and start, and single step forward and backward, to appropriately visualize the rover behavior. Note that this interface is identical whether viewing the results of a simulation or recorded state information from the actual rover making it easy to visually compare the behavior of the rovers.

4. Future Directions

There are several levels of future work to be done on the RSVP tool suite and the processes that utilize them. The lowest level is the addition of new features while keeping the basic paradigm the same. Another level of future work is the extension of the tools into more immersive environments. A third area of interest is in collaboration for building command sequences. The current mission has utilized parallel development only by splitting the sequences logically and having multiple operators work on separate sections. Another area of future work is the support of upcoming missions. One such mission is a new rover under development that is six-sided, has six legs, each terminated with a wheel, and the ability to roll over smooth terrain and walk over rougher terrain. A longer term area of future work is moving up the ladder of onboard autonomy. The MER rovers have commands with some level of autonomy but future systems will have higher levels of autonomy, with goals to achieve, autonomous onboard evaluation of the science value of local targets, interaction of instruments, and other features. These will require a more advanced type of commanding with goal specification and constraint evaluation.

Acknowledgements

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[BIESIADECKI05] J. Biesiadecki *et al.*, "Mars Exploration Rover Surface Operations: Driving Opportunity at Meridiani Planum," Proc. IEEE-SMC, Waikoloa, HI, October 2005.

[COOPER98] B. Cooper, "Driving on the Surface of Mars Using the Rover Control Workstation," Proc. SpaceOps '98, Tokyo, Japan, 1998.

[HARTMAN05] F. Hartman *et al*., "Data Visualization for Effective Rover Sequencing," Proc. IEEE-SMC, Waikoloa, HI, October 2005.

[LEGER05] C. Leger *et al.*, "Mars Exploration Rover Surface Operations: Driving Spirit at Gusev Crater," Proc. IEEE-SMC, Waikoloa, HI, October 2005.

[MAXWELL04] S. Maxwell *et al* "The Design and Architecture of the Mars Exploration Rovers Project's Rover Sequencing and Visualization Program," Proc. SpaceOps '04, Montreal, Canada, 2004.

[MAXWELL05] S. Maxwell *et al* "The Best of Both Worlds: Integrating Textual and Visual Command Interfaces for Mars Rover Operations," Proc. IEEE-SMC, Waikoloa, HI, October 2005.

[WRIGHT98] J. Wright *et al,* "Immersive Environments for Mission Operations: Beyond Mars Pathfinder", Proc. SpaceOps '98, Tokyo, Japan, 1998.

[WRIGHT04] J. Wright *et al,* "Terrain Modelling for Immersive Visualization for the Mars Exploration Rovers", Proc. SpaceOps '04, Montreal, Canada, 2004.

[WRIGHT05] J. Wright *et al,* "Terrain Modelling for In-situ Activity Planning and Rehearsal for the Mars Exploration Rovers," Proc. IEEE-SMC, Waikoloa, HI, October 2005.

[YEN05] J. Yen *et al.*, "Sequence Rehearsal and Validation on Surface Operations of the Mars Exploration Rovers," Proc. IEEE-SMC, Waikoloa, HI, October 2005.