RISK ANALYSIS SIMULATION OF ROVER OPERATIONS FOR MARS SURFACE EXPLORATION

Robert Shishko, Ph.D.
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109 USA

Abstract

Risk management advocates have long sought to directly influence the early stages of the systems engineering process through a more effective role in system design trade studies. The principal obstacle to this has been the lack of credible ways to represent and quantify mission risk—that is, a combination of the probability of mission success ("system safety") and science value—for the project manager and the rest of the design team. If it were possible to quantify mission risk, then the effects of proposed mission and system design changes could be calculated, and along with life-cycle costs, could be used to explore the design space more extensively and select better designs.

JPL has been working to build the capability to quantify the probability of mission success using a federation of diverse simulations and models, each of which contributes some vital piece of the puzzle. The initial institutional focus has been on Mars surface operations. This *ensemble computing* framework enables the diverse models and simulations to work together seamlessly. Recent work at JPL has demonstrated the capability to exercise this ensemble from end-to-end using an Oracle-based database to automatically move results from one model/simulation to the next stage in the analysis.

Introduction

Risk management advocates have long sought to directly influence the early stages of the systems engineering process through a more effective role in system design trade studies. The principal obstacle to this has been the lack of credible ways to represent and quantify mission risk—that is, a combination of the probability of mission success ("system safety") and science value—for the project manager and the rest of the design team. If it were possible to quantify mission risk, then the effects of proposed mission and system design changes could be calculated, and along with life-cycle costs, could be used to explore the design space more extensively and select better designs.[1]

Calculating mission return in a probabilistic way leads to some very natural measures of (system) effectiveness (MoEs) for the mission. These risk-based MoEs can show the project manager (or other decision maker), *for a given design*, what confidence is associated with each level of mission return, or alternatively, what design improvements (or descopes) are needed in order to reach a given level of confidence in a particular level of mission return. Further, calculating a probabilistic mission return is an essential step toward building MoEs that can take into account the project manager's risk aversion—that is, how much the project manager is willing to pay to avoid adverse outcomes on the tail of the mission return probability distribution. Ultimately, these capabilities can provide a means of understanding the tradeoff between the probability of mission success and science value, leading to what one researcher has called "the

[1] When the concept of probabilistic mission return is introduced, one must be careful in defining what is meant by a "better" design. A strong definition involves *stochastic dominance*. Design Alternative A stochastically dominates Alternative B if A's mission return is greater than or equal to B's at each probability level. One might also consider situations in which Alternative A is a little worse than B at near-nominal conditions, but a great deal better when off-nominal conditions are encountered. Alternative A may then be considered a more robust design. Choosing between alternatives in which stochastic dominance does not occur is usually handled by picking the one that maximizes the expected (von Neumann) utility, where the utility function is defined over the domain of mission return.

risk-adjusted mission value" (i.e., the project manager's certainty equivalent taking into account his/her risk aversion).

Technical Approach

JPL has been working to build the capability to quantify the probability of mission success using a federation of diverse simulations and models, each of which contributes some vital piece of the puzzle. The initial institutional focus has been on Mars surface operations. This *ensemble computing* framework enables the diverse models and simulations to work together seamlessly. The framework assembles the following elements:

- Decision tree representation of potential space environments (surface terrain, atmospheric opacity, near-surface temperature cycles) and their probabilities
- Potential space environments are simulated at the level of resolution appropriate for determining interactions of the rover with the environment (scalable synthetic environments)
- Virtual rovers are represented in these synthetic environments and commanded to complete a virtual mission. Use of supercomputer assets allows thousands of Monte Carlo trials to be run, which generate failure-driver pdf's
- Hardware reliability simulation (based on component-level Probabilistic Physics of Failure models) generates a system reliability for each Monte Carlo trial
- Integration of all Monte Carlo trials (via the original decision tree) into a system safety pdf

Recent work at JPL has demonstrated the capability to exercise this ensemble from end-to-end using an Oracle-based database to automatically move results from one model/simulation to the next stage in the analysis. One of the early uses of the ensemble will be to support a navigation software "bake-off" by hosting alternative algorithms on the same rover and mission. Another will be to estimate the rover's ability to study various science sites for the Mars '07 mission.

The structure of the framework is elaborated in Figure 1.

A Collaborative Process for Risk Analysis

The elements of the ensemble computing framework have been joined to a collaborative process for initiating and completing a particular risk analysis since many diverse disciplines must work together to produce, for example, MoEs for a trade study. The players in this process include system and subsystem engineers, scientists, risk managers,[2] simulation specialists, and reliability engineers.

In our collaborative process, the risk engineer/manager plays the key role in moving the process toward completion. An analysis is usually initiated by the project-level system engineer working with a risk engineer/manager. These two establish the hardware and software configuration of the rover(s) to be simulated in the analysis, the Mars site to be explored, and with the help of the mission designers, the landed mission start time. Usually, a performance requirement for the rover's total travel distance during the mission is provided by the project-level system engineer, which may be derived in part from an estimate of the landing dispersion.

Next, the risk engineer works with the project scientist to establish a probabilistic description of the environmental conditions at the selected landing/science site. Together, they use the decision tree method/tool to elicit the scientist's beliefs about surface, atmospheric, and near-surface thermal conditions, and their probabilities. Of course, the scientist uses whatever hard data are available from actual observations in the elicitation process. The results of the elicitation step is a set of quantitative values and concomitant probabilities describing the surface terrain (including relative rock density, average slope, and surface "roughness"), atmospheric opacity (tau values), and diurnal temperature minimums and maximums (including an estimate of their statistical dispersion). Atmospheric opacity and diurnal temperatures are causally connected, so these values and probabilities are elicited conditionally. Figure 2 shows part of the user interface of the decision tree tool used in the elicitation.

[2] The terms risk analyst or risk engineer may be substituted for risk manager.

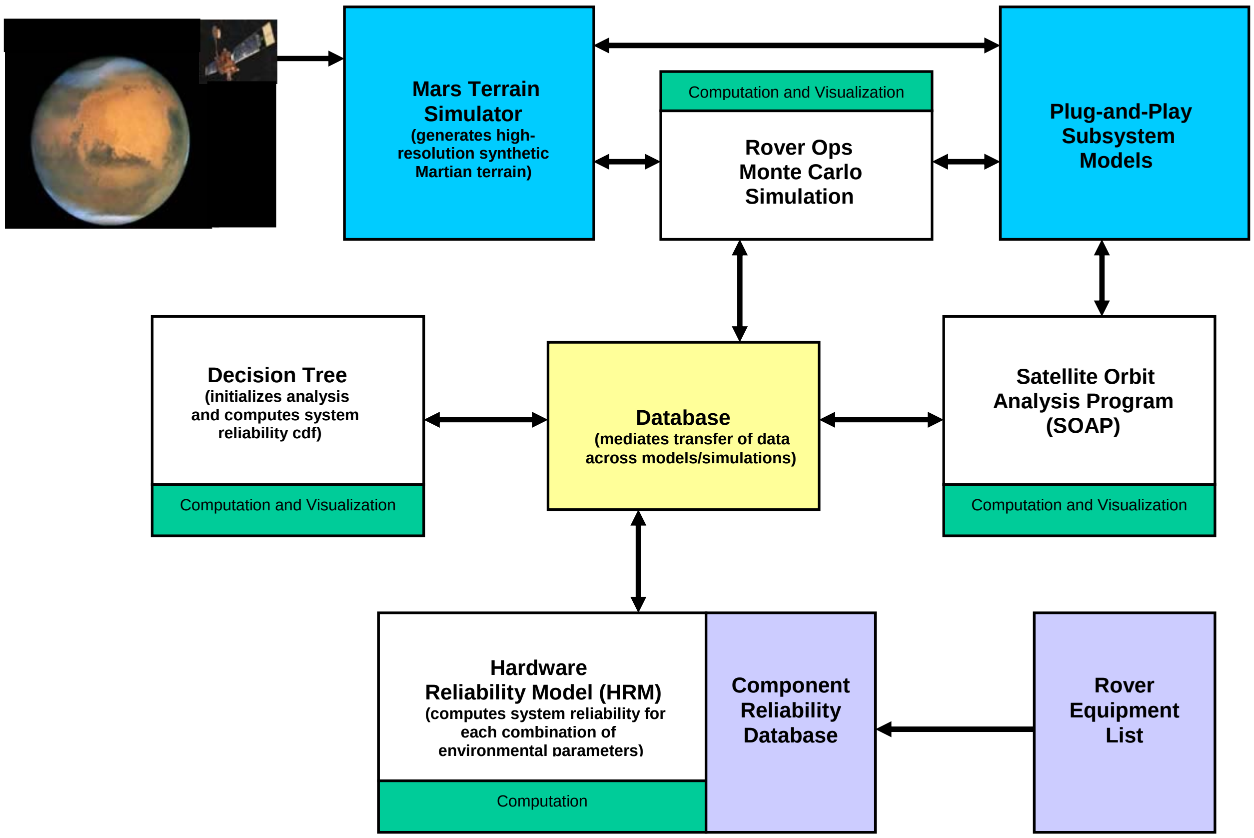

Figure 1—*The risk analysis framework federates a number of models and simulations. Each is described in more detail in the text.*

The risk engineer now has enough information to establish a set of "cases" that need to be simulated to form the analysis. The risk engineer decides at this point how many Monte Carlo trials to request for each case. This is usually based on his/her experience with convergence rates for this type of simulation. Along with the results of the above elicitation, the case information is transferred to the mediating database.

At this point the simulation specialist verifies that the requested cases can be performed and sets up the necessary synthetic environments. This requires accessing a library of preprocessed sites or initiating actions to create new synthetic terrains from a library of images.[3] The simulation specialist also invokes the virtual rover hardware and software configuration called for by the system engineer, and lastly initiates the Monte Carlo simulations for all the cases.

The results of each trial within each case are exported to the mediating database during the simulation. Typically, the statistics gathered can be used to generate a pdf of each potential failure driver, such as time, distance traveled, and on-off cycles. When all trials have been completed, the focus of the analysis shifts to the Hardware Reliability Model (HRM). For each trial, the HRM computes the likelihood of no critical failures (system reliability). These data are also placed into the mediating database.

[3] Algorithms for performing this step start with known topographical features from orbital images and construct a statistical and geological correct synthetic terrain. [1-2] The elicited terrain parameters are fed to these algorithms in the Mars Terrain Simulator through the mediating database.

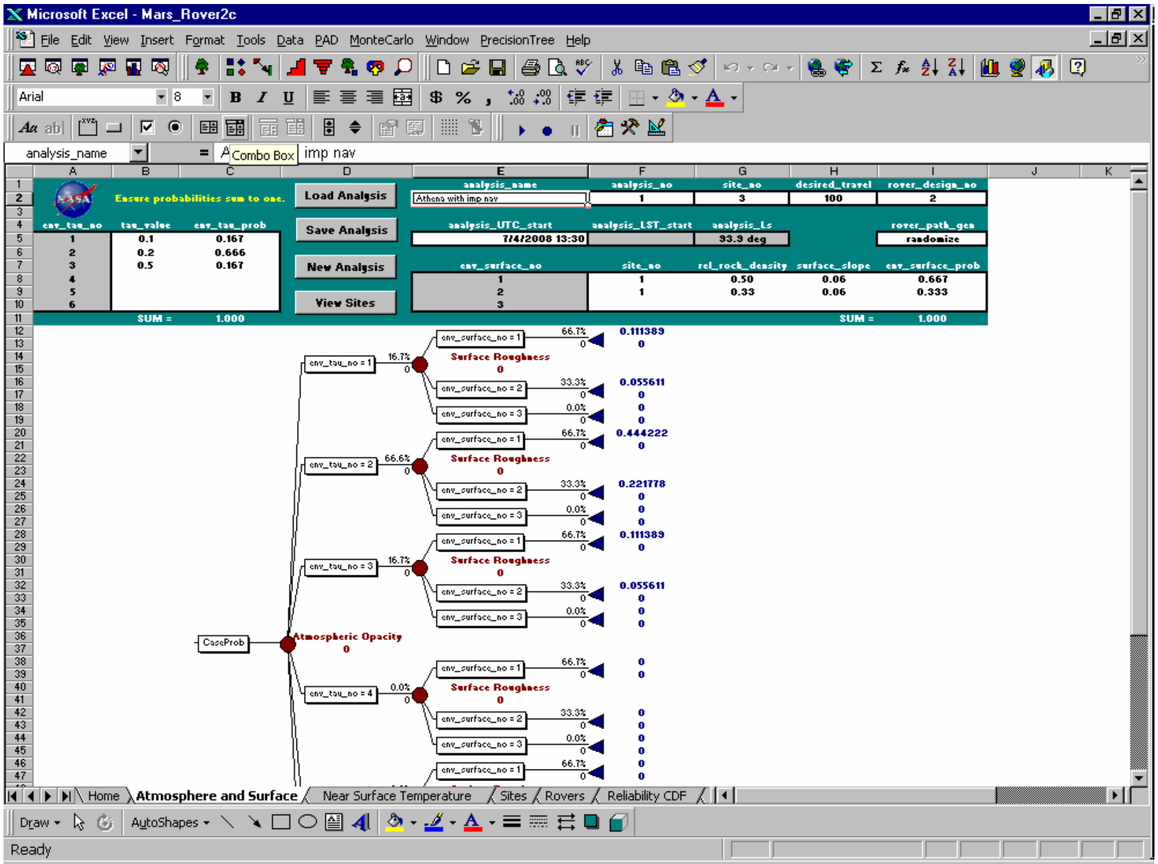

Figure 2—*Working user interface for eliciting information about the surface and atmosphere at a Mars science site. Data are transferred to a mediating database so that other models/simulations have access to them.*

The HRM bases its calculations on data provided by the simulation on failure drivers and on component-level reliability data provided by the reliability engineer. The HRM user interface serves mainly to input these component-level data. The reliability engineer has the job of determining the dominant failure modes within the context of the mission. The reliability engineer then enters data into a normalized relational reliability database that can also serve as a data source for other in-situ rover PRAs. The first of two primary tables of the database defines the relationship between a component and the rover configuration containing it, while the second defines the failure modes for each component, the associated failure driver, failure density functional form (exponential, lognormal, Weibull, etc.), and parameters of the distribution.

The collection of these data is a significant activity for the reliability engineer, since quantitative data of this sort are very scarce. Our HRM database contains rover data from *Sojourner* component testing and Probabilistic Physics of Failure (PPOF) modeling. Eventually Mars Exploration Rover (MER) data will be incorporated.

The focus of the analysis now shifts back to the risk engineer, who uses the original decision tree method/tool to complete the analysis. The tool reads in the reliability results from the mediating database, and automatically creates the system reliability cdf for that analysis by merging system reliabilities with the probability estimates for each combination of Mars environmental parameters. The tool then plots the cdf on a graph. The risk engineer can also automatically include the cdf graph from the previous analysis so that the two alternatives may be compared. The decision tree tool then exports the current cdf's mean, median (50^{th} percentile), and 10^{th} percentile value[4] to the

[4] The 10^{th} percentile system reliability is the value such that the probability that the actual value is higher equals 90 percent.

database. These values can serve as risk-based system MoEs.

If the project manager and/or project-level system engineer believe that the risk-based MoE is unacceptably low, he/she can:

- Change the rover design by adding redundancy, or raise subsystem reliability and/or performance requirements;

- Improve the precision landing capability to reduce the likely travel distance to the science site;

- Change the science site to increase the probability of a "smoother" one, or change the areocentric[5] longitude to reduce the probability of an adverse optical depth; or

- Some combination of the above.

Each of these changes affects the risk-based MoE, but has (science value, mass, power, etc.) implications for the entire project that must be understood. Each alternative proposed change must also be fully costed before one is chosen.

However, before this Monte Carlo simulation-based approach to Mars surface operations risk analysis is practical within a real NASA mission, the turnaround time must be a few hours, or at worst, overnight. Consequently, we have devoted considerable effort to reducing simulation execution time. The architecture in Figure 1 has been implemented using Caltech/JPL supercomputer resources. This was highly desirable because of the number of cases (sets of environmental parameters) and trials (to ensure usable statistical characterizations) that typically need to be run.

The supercomputer implements a "one processor — one rover ("trial") approach, allowing a two orders of magnitude improvement over earlier implementations; but beyond that, a third order of magnitude improvement was achieved by creating a "terrain server", which feeds small patches of high-resolution synthetic Martian terrain (from a much larger swath) to each rover as it moves towards its goal.

The simultaneous movement of many rovers can be rather dramatically displayed while the simulation is executing on the supercomputer. This visualization capability is more than just entertainment. It can be very useful in understanding the behavior of the rover under a variety of conditions that could not be created during system testing.[6]

Figure 3 shows 32 *Sojourner* rovers attempting to reach a common point equidistant from each. The simulation used the actual flight software to control the rover's movement over the synthetic Martian terrain. The visualization records this movement so that the path each took appears as "tracks".[7] The visualization can be captured as a file, and replayed at a future date.

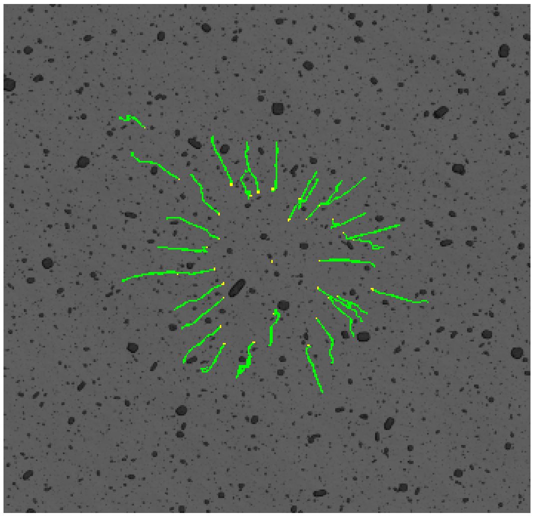

Figure 3—*This visualization was captured from a rover ops simulation involving 32 rovers moving simultaneously over synthetically generated 33 percent Viking Lander 2 (VL2) terrain—that is, terrain conforming the VL2 rock size frequency distribution, but with a rock density covering just 5.9 percent of the terrain instead of 17.6 percent.*

[5] Time of year on Mars is usually denoted by areocentric longitude, where zero degrees refers to the vernal equinox in the northern hemisphere.

[6] For example, the visualization showed that *Sojourner's* navigation algorithm had a propensity for inducing "confused" behavior, which helped the algorithm designer improve the next generation of software once the behavior was identified.

[7] Of course, a more sophisticated randomization method (than initializing the rovers' positions on a circle) is used in the simulation.

For missions of substantial duration, the rover ops simulation needs information regarding its relationship to other objects—for example, the Earth, a potentially interesting Mars science site, and any Mars orbiters with which it might communicate. For these physics-based time-dependent data items, the risk analysis framework incorporates the Satellite Orbit Analysis Program (SOAP). SOAP was developed by the Aerospace Corporation originally for visualization of Earth orbiters, but has been substantially modified to serve as a visualization and computation engine for Mars missions. [3]

Specifically, SOAP has been modified to provide (a) view periods for in-place and potential Mars communication satellites, and the Earth-based DSN; (b) local solar time and sun angle; and (c) an interface to the rover ops simulation through which the two simulations can be synchronized. Figure 4 shows SOAP simulation output that might support a Mars '07 mission.

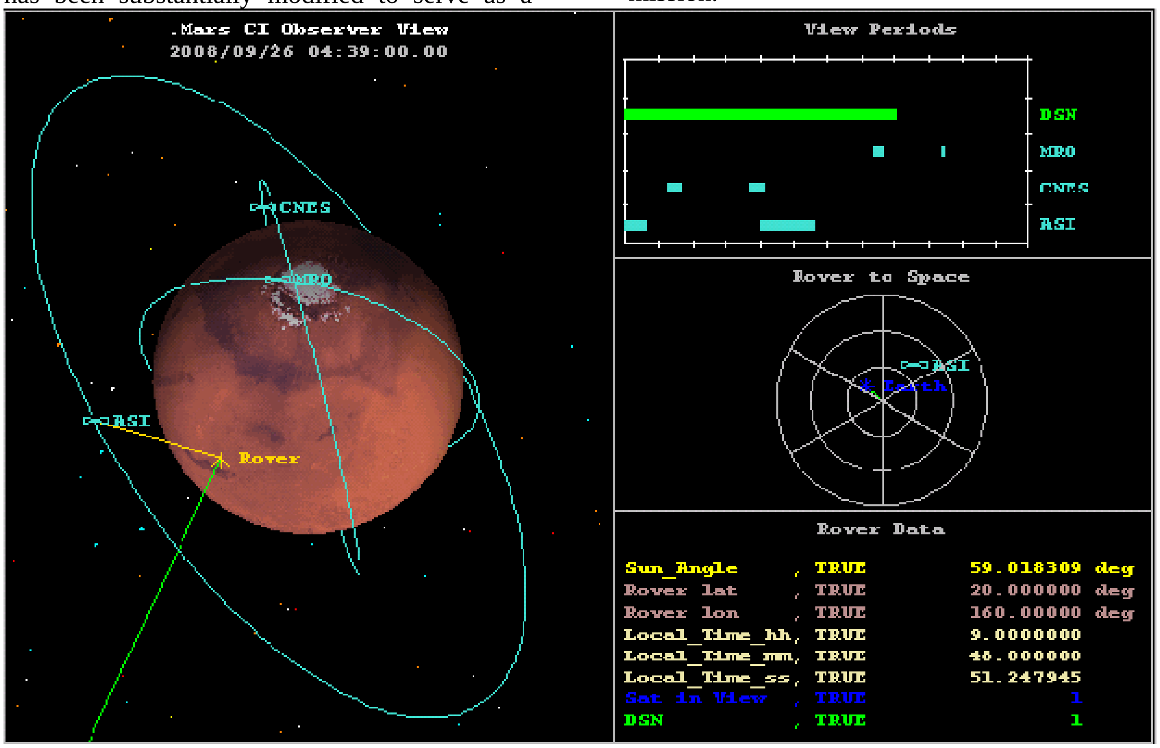

Figure 4—*This SOAP screen capture shows some of the data the SOAP calculates and can be used by the rover ops simulation. The visualization capabilities aid in understanding interactions.*

Results for *Sojourner*

The first complete risk analysis of this type was performed for the Mars Pathfinder (MPF) *Sojourner* in 1997 before the current ensemble computing framework was developed. It provided insights into the issues that we needed to address to get to the current framework. [4] The *Sojourner* surface operations simulation trials were executed serially and data were "passed" from one model/simulation as ASCII files.

Some results of two separate surface operations simulations are shown in Figure 5. The figure shows the Weibull probability density functions for the distance the *Sojourner* would have to travel in order to reach a target 100 meters away (geodesic distance). In the simulations, the actual *Sojourner* flight software controlled the movement of the rover over two synthetic, but high (i.e., centimeter-level) resolution, Martian terrains. The right-most curve resulted from terrain that we describe as a 50 percent Viking Lander 2 (VL2) site. By that we mean that the statistical size-frequency distribution of rocks and craters was the same as Golembek's characterization for VL2 [5-7], except that the absolute number of rocks and craters (centers per square meter) was only half of the actual VL2 site. The left-most curve resulted from terrain that we describe as a 25 percent VL2 site.

The distance traveled is described by a probability density because each simulation was actually rerun many times randomizing the initial location of the rover relative to the target. The individual simulation trials were converted into a Weibull probability density function using well-known parameter estimation techniques. [8] The simulations, shown in Figure 5, were run under conditions similar to the actual MPF landing site—high optical depth (very clear), 19.5 degrees N latitude, and 143 degrees areocentric longitude (mid-to-late summer).

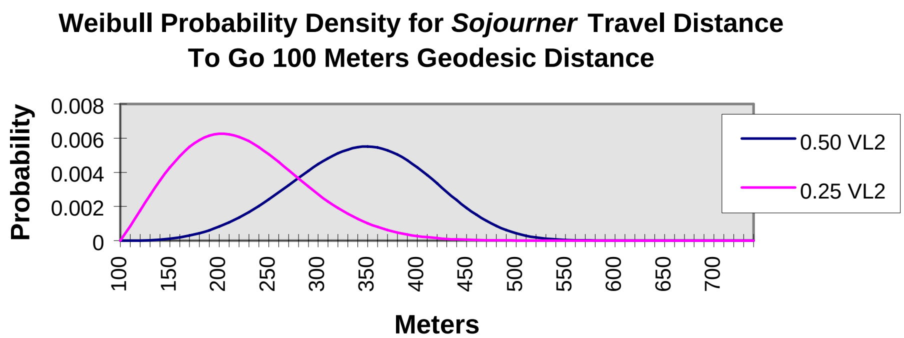

Figure 5—*Some Results of the Sojourner Surface Operations Simulations. Each curve is derived from individual Monte Carlo trials for that particular case. The simulation results from each trial are used by the Hardware Reliability Model to compute the system-level reliability for that trial.*

One role of the surface operations simulations was to provide quantitative values for the failure drivers (e.g., operating time, distance driven, on-off cycles) in the individual rover component failure models. One of these models was for the rubbing rotor, which is found in each of the six rover wheels. If the *Sojourner* had to travel the expected distance of 343 meters at a 50 percent VL2 site to reach a target 100 meters away, then the reliability of the rubbing rotors was calculated at 0.9489 using the rubbing rotor failure model (for the nominal diurnal temperature cycle) and recognizing that only five out of six had to work. The HRM aggregated all of these individual component results to create an overall system reliability for a particular combination of Martian environmental parameters (including variations in the diurnal temperature cycles). The system reliability results are rolled up over the set of Martian environments described in the decision tree model to produce a probabilistic description of the rover's reliability in traveling 100 meters geodesic distance. This cdf is displayed in Figure 6.

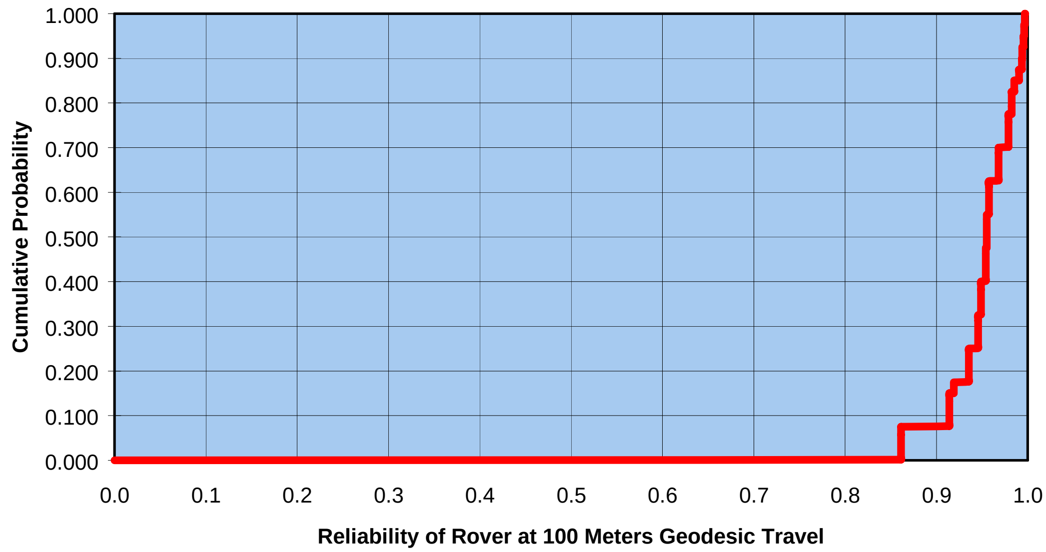

Figure 6—*Risk-Based Measure of Effectiveness for Sojourner Operations. This curve resulted from the Mars environmental parameter values and their respective probabilities represented in the decision tree, the failure driver values computed in the rover operations simulations, and the reliability characteristics of the rover components.*

The curve results from the uncertainties in the Martian environment at the landing site and the consequent effects the environment has on rover operations and reliability. The mean (i.e., expected value of) system reliability, which in this case represents a good choice for risk tracking during development, is 0.952.

Coincidentally, the choice of 100 meters for the surface operations simulation came remarkably close to the actual figure of 104 meters for the *Sojourner's* total geodesic travel distance during the 83 sols of operations. The actual distance traveled during the same time was 150 meters as measured by the wheel odometry—well within a standard deviation from the mean travel distance of the 25 percent VL2 terrain. Unfortunately, the primary validation of the surface operations simulation, during which *Sojourner* was to traverse about 40 meters of terrain autonomously, was never performed. The instructions were in the sequence to be uplinked when contact with the MPF lander was lost on September 27, 1997.

The *Sojourner* simulations, however, strongly confirmed that the rock density at the MPF landing site was a legitimate concern. The simulations showed that the effects of rock size-frequency distributions on rover travel distance, and hence reliability, are nonlinear

Conclusions

Probabilistic mission simulation has a great deal of potential for future rover missions such as Mars '07 and the Mars Sample Return because of the sensitivity of the results to rover design, landing site, and other mission parameters. It becomes possible to use "risk as a resource" in a concrete way during trade studies. Further, there appears to be no known obstacle to performing advanced risk analyses using this approach. Ultimately, the results of rover risk analyses of this type are likely to have important implications for future rover design (reliability, autonomy, etc.) and overall risk mitigation strategies (mission redundancy, hedges, precision landing accuracy improvements).

Improvements in the current implementation of the framework, however, are needed. An open architecture with a true plug-and-play capability for rover subsystems models has yet to be implemented. This capability would speed the process and lower the cost of capturing alternative designs. It is not unreasonable to want to test new rover navigation algorithms, sensor types, power, and mobility subsystems without having to overhaul the rover operations simulation.

Another area for improvement is in the way extended operations are simulated; whereas the *Sojourner* operations simulation covered 100 meters in one to 10 sols, future simulations will have to deal with traverses of 10 to 50 km over a period of 10 to 100 sols.

References

[1] Gaskell, RW 1993, Martian Surface Simulations, *Journal of Geophysical Research*, v98 (E6), 11099-11103

[2] Gaskell, RW, Husman, LE, et al. 2001, Synthetic Environments for Simulated Missions, IEEE Paper 13.0406, *Proceedings of the IEEE Aerospace Conference*, March 2001, Big Sky, MT. IEEE 0-7803-6599-2/01

[3] Stodden D, et al. 2001, *SOAP Version 11 User's Manual*, Aerospace Corporation

[4] Shishko, R 1997, Risk Analysis and System Trades in the Mars Sample Return Mission, *Proceedings of the 1st Joint ESA/INCOSE Conference on Systems Engineering*, November 1997, Noordwijk, NE. WPP-130

[5] Golombek M, and D Rapp 1997, Size-Frequency Distribution of Rocks on Mars and Earth Analog Sites: Implications for Future Landed Missions, *Journal of Geophysical Research—Planets*, v102 (E2), 4117-4129

[6] Haldemann, AFC, et al. 2000, Rock Statistics at the Mars Pathfinder Landing Site, Roughness and Roving on Mars, *Workshop on Mars 2001*, 2527.pdf

[7] Bernard, DE and Golembek MP 2001, Crater and Rock Hazard Modeling for Mars Landing, *Proceedings of the AIAA Space 2001 Conference*, Albuquerque, NM, August 2001. AIAA-2001-4697

[8] McCormick, NJ 1981, *Reliability and Risk Analysis*, Academic Press, Inc. Orlando

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