2007

Planning to Fail - Reliability as a Design Parameter for Planetary Rover Missions

S. Stancliff
Carnegie Mellon University

John Dolan
Carnegie Mellon University

A. Trebi-Ollennu
NASA Jet Propulsion Laboratory

Planning to Fail - Reliability as a Design Parameter for Planetary Rover Missions

S. Stancliff, J. Dolan
The Robotics Institute
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA, USA
{sbs,jmd}@cs.cmu.edu

A. Trebi-Ollennu
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA, USA
Ashitey.Trebi-Ollennu@jpl.nasa.gov

***Abstract*—**The Mars Exploration Rovers (MER) have been operating on Mars for more than three years. The extremely high reliability demonstrated by these rovers is a great success story in robotic design. This reliability comes at a high cost, however, both in the initial cost of developing the rovers and in the ongoing operational costs for their mission extensions. If it were possible to design rovers with reliability more in line with their mission requirements (in the case of MER, 90 days), considerable cost reductions could be achieved. This will be even more important for future planetary robotic missions due to greatly increased mission durations.

In this paper we present an overview of our ongoing research in the area of predicting robot mission reliability, and we show how a mission designer can trade off reliability against costs in order to find an optimal reliability target for a given robotic mission. Our results show that for a given mission there is an optimal reliability range with respect to cost and that having rovers with reliability that is too low or too high is suboptimal from an economic standpoint. This suggests that a better cost-reliability tradeoff can be obtained by "planning to fail" by designing rovers which have lower reliability than current legacy designs.

***Keywords*:** *planetary rovers, mission design, mission cost, reliability, failure, risk.*

I. Introduction

In the near future, NASA intends to send rovers to Mars for missions lasting an order of magnitude longer than the intended duration of the Mars Exploration Rovers (MER) mission. If these future rovers follow legacy designs, then increasing the mission duration by an order of magnitude will require that the rovers be built using components with failure rates an order of magnitude lower. Since NASA rovers already make use of some of the most reliable components available, it is doubtful whether components with order of magnitude lower failure rates are available, let alone affordable.

In order to increase rover mission durations without incurring exponential increases in rover costs, it is necessary to consider risk not simply as something to be minimized to the greatest extent possible, but instead as a quantitative design factor to be traded off against other design factors in order to seek an optimal mission configuration.

In the mobile robotics literature there is little formal discussion of reliability and failure. When reliability is mentioned, it is usually qualitatively, and in passing. Reference [1], for example, mentions intermittent hardware failures as an explanation for gaps in experimental data but makes no attempt at characterizing the failures.

A handful of prior papers make use of reliability engineering for analysis of mobile robot failure rates. Reference [2] provides an overview of robot failure rates at the system level (i.e., robot model X failed Y times in Z hours of operation) and also breaks down failures according to the subsystem that failed (actuators, control system, power, and communications). Reference [3] extends the work in [2] both by the inclusion of additional failure data of the same type and also by addition of new categories of failure—those due to human error. Reference [4] provides a detailed analysis of failures experienced by some of the robots used in searching the World Trade Center wreckage in 2001. Reference [5] provides failure data for robots used in long-term experiments as museum guides.

What these papers have in common is that they use reliability engineering tools in the assessment of existing robots. Our work differs in that it addresses how to use reliability engineering tools for designing robots and robotic missions.

In earlier work we have developed methods for using reliability engineering tools to predict the probability of a robot failing during a mission [6], and we have used these tools to compare the performance of different robot and robot team configurations [7]. The only known work preceding ours in the area of predicting mobile robot team reliability is [8]. That paper's methods are similar to ours in that they are based in the reliability literature, but that work has a narrow focus on teams of robots with cannibalistic repair capability. In contrast, we are developing a general methodology that can be applied to a wide variety of robot teams and missions.

The work presented in this paper differs from our earlier work by addressing the relationship between robot reliability and overall mission cost, and demonstrating how this relationship can be used to identify an optimal reliability level which minimizes mission cost.

II. Example mission scenario

A. Missions and Tasks

Consider a planetary exploration mission where a team of rovers is tasked to install a solar panel array for a measurement and observation outpost. The mission consists of carrying 50 solar panels from the landing site to the outpost and then assembling them. The size of the solar panels is such that two rovers are needed to carry and assemble one panel.

For the purposes of the reliability analysis, the task of assembling a solar panel is broken down into three subtasks:

- Transit to the outpost,
- Assemble the panel, and
- Return to the landing site.

We assume that failure occurs only at the end of a subtask. This allows us to avoid dealing with partially completed subtasks. This simplification does not limit the resolution of the representation because tasks can be restated into smaller subtasks if needed.

B. Rovers and Components

For this analysis we assume that the rovers on the team are identical. The rovers are considered to be made up of several subsystems that are independent from the standpoint of reliability. The specific partitioning is not important to the methodology, but for the analyses in this paper the rovers are divided into the subsystems listed in Table 1.

The subsystem reliabilities listed in Table 1 were calculated from the failure rates of the major components in each subsystem. An example component breakdown for the power module is shown in Table 2. Due to the limited amount of failure data available for planetary rovers, the failure rates in Table 2 were derived from the RAC databooks ([9]) which are commonly used for reliability prediction in aerospace and military applications. Additional details on the calculation of subsystem failure and the combining of component failure rates can be found in [10].

We assume that the failure of any single subsystem leads to failure of the entire rover. For the current example mission this is a reasonable assumption, since all of the subsystems must be functioning in order to complete the mission subtasks.

The probability of a subsystem failing during a task is found using standard reliability engineering methods assuming a constant failure rate. Two inputs determine the module failure probability: the module failure rate and the length of time for which the module is operated during the task. The durations shown in Table 3 were assigned using reasonable assumptions about the relative durations of different tasks and the relative usage of different modules. During the transit task, the panels are assumed to be locked in a fixed position not requiring manipulator actuation.

The probability of survival for a subsystem for a given task is given by the equation

$$P = e^{-t\lambda}, \qquad (1)$$

where t is the amount of time that the subsystem is used during the task and λ is the failure rate for the subsystem.

Using (1) and the data from Tables 1 and 3, we calculated the probability that each subsystem will survive each task. These probabilities are shown in Table 4.

TABLE 1
Rover Subsystems and Reliabilities

Subsystem	MTTF (h)
Power	4202
Computation&Sensing	4769
Mobility	19724
Communications	11876
Manipulator	13793

TABLE 2
Components Comprising Power Subsystem

Component	Quantity	Failure Rate (1/h)
Battery	2	2.10×10^{-7}
Battery control board	2	4.00×10^{-7}
Mission clock	1	1.00×10^{-7}
Power distribution unit	1	1.70×10^{-6}
Power control unit	1	1.90×10^{-7}
Shunt limiter	1	1.14×10^{-5}
Electrical heater	2	3.00×10^{-6}
Radioisotope heater	2	1.36×10^{-5}
Thermal switch	2	9.50×10^{-5}

TABLE 3
Subsystem Usage by Task in Hours

Subsystem	Transit	Assemble	Return
Power	6	8	6
Computation&Sensing	6	4	6
Mobility	6	8	6
Communications	2	4	2
Manipulator	0	8	0

III. Approach

The experiments in this paper make use of the method described in [5] for predicting probability of mission completion. In this method, the mission is represented using a state machine that is simulated stochastically.

The simulation is repeated many times, with the average score of all trials giving the overall probability of mission completion. The results of the simulations were verified by hand calculation for a few simple cases.

TABLE 4
Subsystem Probability of Survival by Task

Subsystem	Transit	Assemble	Return
Power	99.86%	99.81%	99.86%
Computation&Sensing	99.87%	99.92%	99.87%
Mobility	99.97%	99.96%	99.97%
Communications	99.98%	99.97%	99.98%
Manipulator	100%	99.94%	100%

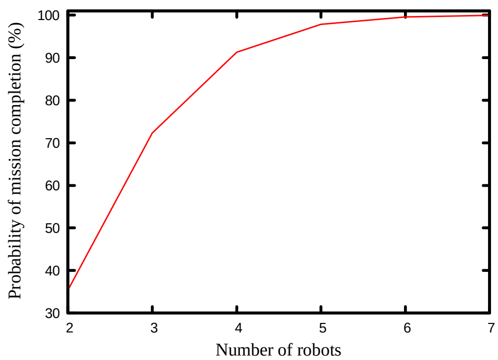

Fig. 1. Mission Reliability as a Function of Team Size

A. Relationship Between Team Size and Mission Success

Using this method, we first examine the relationship between the number of rovers on the team and the probability of completing the mission. Figure 1 compares teams of two to six rovers, each composed of the baseline components with the subsystem reliabilities shown in Table 1. This analysis can be used to determine a minimal team size for a required probability of mission success. For instance, if we set the required probability of mission success at 99.5% then Figure 1 shows that the team must consist of at least six rovers.

B. Relationship Between Component Reliability and Mission Success

Figure 1 shows that a six-rover team exceeds the mission reliability requirement. In such a case, a mission designer may wish to choose lower-reliability components in order to decrease mission costs. The same simulations used to create Figure 1 can be used to determine the minimum component reliabilities required to meet a particular mission reliability requirement. Figure 2 compares six-rover teams using components with reliabilities which vary from 60% to 100% of the values in Table 1. From this we find that we can achieve the 99.5% goal by using components with 95% of the reliabilities shown in Table 1.

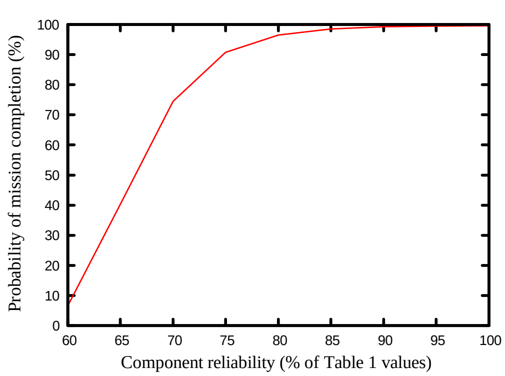

Fig. 2. Mission Reliability as a Function of Component Reliability

We have now determined that the smallest team with the lowest-reliability components which can achieve the design goal of 99.5% probability of mission success is a six rover team with the component reliabilities shown in Table 5. We use this team as the baseline for the comparisons that follow.

C. Relationship Between Component Reliability and Cost

The reliability of the rovers is related to the overall mission cost in two ways. First, there is the increased cost associated with higher-reliability rovers. Second, there is the increased expected value of the mission when using higher-reliability rovers due to a higher probability of mission success.

1) Cost of Reliability: In choosing components from which to build rovers, a designer would usually make choices among a small number of alternative components, each providing a certain reliability for a certain cost. However, in the early stages of design the mission designer may not have complete information about available components. In this case, it is useful to have a parametric model of the cost–reliability relationship. Reference [27] provides a general model for this relationship, which is given as

$$c = \exp\left\{(1-f)\cdot\frac{(R_i - R_{\min})}{(R_{\max} - R_i)}\right\} \quad (2)$$

where R_i is a reliability of interest between R_{min} and R_{max} ; c is the relative cost of R_i compared to R_{min} ; f is the feasibility of reliability improvement (a number between 0 and 1); and c is the resultant relative cost of R_i with respect to R_{min}.

This equation can be used to calculate the relative cost of the components used by the six-rover teams with differing component reliabilities. These costs are plotted in Figure 3 as a percentage of the baseline team cost, using R_{min}=0, R_{max}=1 and f=0.95. We examine the effect of changing the feasibility constant later in this paper.

Launch costs are also affected by rover reliability. More-reliable rovers will weigh more, due to increased size of more-reliable components and due to increased component redundancy. We have not found a model for the reliability–weight relationship in the literature. As an initial approximation for launch costs we assume that the relationship between weight and reliability is directly linear and that the relationship between launch costs and weight is also directly linear.

TABLE 5
COMPONENT RELIABILITIES GIVING 99.5% PROBABILITY OF SUCCESS FOR SIX-ROBOT TEAM

Subsystem	**MTTF (h)**
Power	3992
Computation&Sensing	4531
Mobility	18738
Communications	11282
Manipulator	13103

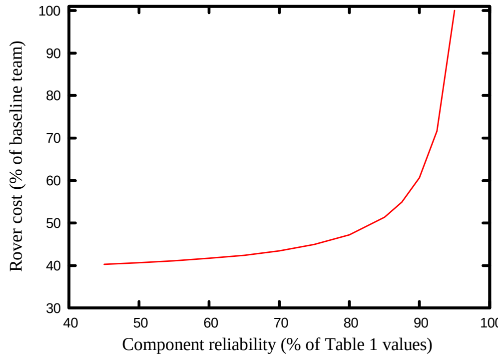

Fig. 3. Relative Cost of Rovers as a Function of Component Reliability

2) Expected Value of Mission: Any mission must have some inherent value to it. For some missions there will be an obvious economic or strategic value to which a dollar amount can be assigned. For a mission that lacks such an obvious dollar value, the cost of the baseline mission itself can be used as a lower bound for this inherent mission value, since the sponsor presumably expects some return on the investment.

Multiplying the probability of mission success by the inherent value of the mission gives an expected value for a given team configuration. For example, the relationship between component reliability and expected mission value is given by Figure 2, with the vertical axis relabeled as "expected value as percent of inherent value".

D. Overall Mission Cost–Reliability Relationship

Taking the expected mission value calculated above and subtracting the rover development and launch costs gives us an estimate of the net expected gain for the mission. We ignore operating costs here since we expect them to be roughly constant with respect to rover reliability (probably slightly higher for lower-reliability rovers due to the increased need for intervention).

In order to combine these costs meaningfully, we assign real dollar values to the various costs for the baseline team. These values are estimated from the costs of the MER mission, along with the assumption that the rovers for this mission would be somewhat cheaper and smaller than the MER rovers due to advances in technology and also because they are single-purpose machines. The values we assigned for the baseline team are shown in Table 6. Figure 4 then plots these component costs and values as well as the net expected gain as a function of rover component reliability.

IV. Conclusions

The most significant thing revealed by Figure 4 is that there is clearly an optimal reliability range with respect to the expected gain of the mission, and that this optimal reliability is significantly lower than the reliability of the baseline (legacy) design.

TABLE 6
BASELINE TEAM COSTS AND REWARDS

Item	Cost ($ Millions)
Robot cost (entire team)	150
Launch cost (entire team)	300
Inherent value of mission	450

The shape of the expected gain curve shows that for low-reliability rovers the cost of failure drives the expected gain value down, while for very high-reliability rovers the high cost of the rovers themselves drives the expected gain down. The optimal reliability range therefore lies in a medium-reliability region where neither of these costs is as high.

In order to evaluate the effects of some of our assumptions on these conclusions, we have repeated the above analysis for different values of the feasibility constant (since this value was arbitrary) and of the mission inherent value (since we used a lower-bound estimate for this value). These results are shown in Figures 5 and 6. These figures show that while the shape of the expected gain curve changes somewhat with these parameters, the overall trends remain the same, and both figures support the argument that the optimal range for mission reliability is at a lower level than we would intuitively consider to be the case.

While we expect that these curves will vary for different missions, we expect that the general trends will hold, indicating that it can be economically wiser to "plan to fail" by building rovers which have lower reliability than current legacy designs.

Acknowledgement

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration and funded through the Director's Research and Development Fund Program.

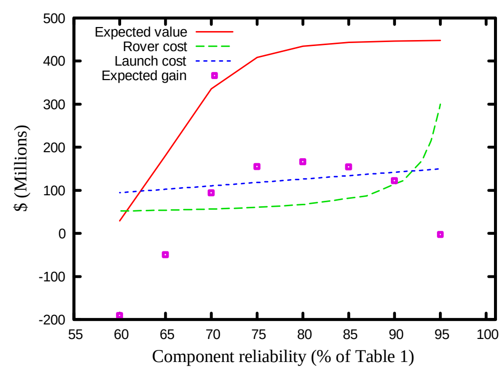

Fig. 4. Net Expected Gain with f=0.95, value = $450M

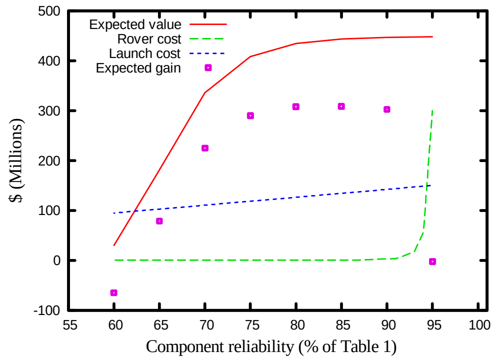

Fig. 5. Net Expected Gain with f=0.5, value=$450M

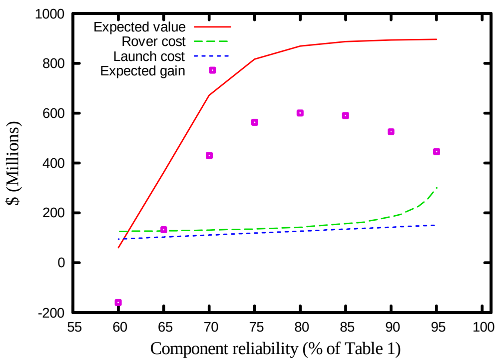

Fig. 6. Net Expected Gain with f=0.95, value= $900M

REFERENCES

[1] R. Gockley, J. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A. Schultz and J. Wang, "Designing robots for long-term social interaction," *Proc. 2005 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS 2005)*, pp. 1338–1343.

[2] J. Carlson and R. Murphy, "Reliability analysis of mobile robots," *Proc. 2003 IEEE Int. Conf. Robotics and Automation (ICRA 2003)*, pp. 274–281.

[3] J. Carlson, R. Murphy and A. Nelson, "Follow-up analysis of mobile robot failures," *Proc. 2004 IEEE Int. Conf. Robotics and Automation (ICRA 2004)*, pp. 4987–4994.

[4] M. Micire, "Analysis of the robotic-assisted search and rescue response to the World Trade Center disaster," M.S. thesis, University of South Florida, 2002.

[5] I. Nourbakhsh, C. Kunz and T. Willeke, "The mobot museum robot installations: a five year experiment," *Proc. 2003 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS 2003)*, pp. 3636–3641.

[6] S. Stancliff, J. Dolan, and A. Trebi-Ollennu, "Mission reliability estimation for repairable robot teams," *Int. J. Advanced Robotic Systems*, Vol. 3, No. 2, Jun., 2006, pp. 155–164.

[7] S. Stancliff, J. Dolan, and A. Trebi-Ollennu, "Mission reliability estimation for multirobot team design," *Proc. 2006 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2006)*, pp. 2206–2211.

[8] C. Bererton and P. Khosla, "An analysis of cooperative repair Capabilities in a team of robots," *Proc. 2002 IEEE Int. Conf. Robotics and Automation (ICRA 2002)*, pp. 476–486.

[9] RAC Automated Databook, Version 2.20, IIT Research Institute, 1999.

[10] S. Stancliff, J. Dolan, and A. Trebi-Ollennu, "Towards a predictive model of mobile robot reliability," Carnegie Mellon University technical report CMU-RI-TR-05-38, 2005.