2006

Mission Reliability Estimation for Multirobot Team Design

S. B. Stancliff
Carnegie Mellon University

John M. Dolan
Carnegie Mellon University

A. Trebi-Ollennu
NASA Jet Propulsion Laboratory

Published In
.

Mission Reliability Estimation for Multirobot Team Design

S.B. Stancliff and J.M. Dolan
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213 USA
stancliff@cmu.edu, jmd@cs.cmu.edu

A. Trebi-Ollennu
Jet Propulsion Laboratory
Pasadena, CA 91109 USA
Ashitey.Trebi-Ollennu@jpl.nasa.gov

***Abstract* — One reason given for the use of multirobot systems is that many cheap robots are more reliable than one expensive robot. To date, however, there has been no quantitative analysis to support this assertion. This paper presents the first quantitative support for the argument that larger teams of less-reliable robots can perform certain missions more reliably than smaller teams of more-reliable robots. Our results show that for short missions, in fact, a team of four robots can provide greater mission reliability than a team of two robots, even when the individual robots in the team of four have reliability that is an order of magnitude lower. These results suggest that considerable cost reductions can be achieved for some missions by choosing larger teams of less-reliable robots over smaller teams of more-reliable robots.**

Index Terms* — Mobile robots, multirobot systems, mission design, reliability.

I. Introduction

Applications of multirobot systems can be divided into two categories: those where multiple robots are necessary for task completion, and those where a single robot could complete the task but where multiple robots are desirable for reasons beyond task completion. An example task falling into the first category is soccer—a single robot cannot play soccer. An example task in the second category is area coverage—while in many cases an area can be covered by a single robot, it may nonetheless be preferable to use more than one robot.

When the mission itself does not dictate a particular robot team configuration, there are multiple requirements which a mission designer must consider. Three important factors which we consider here are time, cost, and reliability.

Time can be a reason for using more robots than the minimum required because, for some tasks, having extra robots can reduce the time required to complete the task. For instance, in an area coverage task, multiple robots can work in parallel in order to accomplish the task more quickly.

Cost is an important consideration in team size. There is the cost of additional robots. There is the cost of robot components—more robust components cost more. There are operating costs such as transportation and maintenance, which may be higher for a larger team. Infrastructure costs may be greater for a larger team; for instance, a larger team may require more communications bandwidth.

The third performance criterion we consider here is reliability, expressed as the probability of mission completion (PoMC). A requirement for a mission to have a certain probability of successful completion can dictate the minimum number of robots required for the mission. For example, if one robot has a 90% probability of surviving a task, but the mission requirement is for a 97% probability of having one robot survive the task, then one way to meet this requirement is by sending two robots (giving a 99% chance that one would survive).

These criteria (time, cost, reliability) are highly interdependent. As an example, adding more robots to a mission increases the cost, but it can also reduce the amount of time required to complete the mission. Reducing the mission duration means that the robots don't need to survive as long, so they can be built of lower-reliability components, which reduces the cost.

These relationships among team size, component reliability, cost, time, and mission success have been mentioned in the robotics literature, but only in passing and only in qualitative terms. In particular, researchers often claim that multirobot systems provide greater reliability than single-robot systems (e.g., [1,2,3,4]).

Superficially, such a claim seems obviously true—if three robots are sent to do a task instead of one, there is a greater chance of completing the task. When one examines the above claim in greater depth, however, finding the answer can be complicated. In this example, the cost of completing the task has been tripled by sending three robots. If these same additional funds were instead invested towards improving the reliability of a single robot, then which would be more likely to complete the task—the three robots or the single superior robot? The answer is no longer obvious.

In this paper, we provide quantitative analyses of the tradeoffs among these design variables. For a sample robot mission, we compare the reliability and cost of teams with differing numbers of robots and different robot reliabilities. We examine questions which a mission designer would want to ask, such as "For a given mission and number of robots, what is the minimum robot reliability required to provide a certain probability of mission completion?" and "For a given mission, if I use extra robots, how much less reliable can they be and still give the same probability of completing the mission?"

This paper makes use of the methodology we have developed previously for predicting the reliability of robot teams (see [5]). While the reliability engineering literature provides methods for predicting the reliability of systems composed of independent components, the nature of multirobot systems is such that there is a significant amount of dependence among the reliabilities of team members. In [5] we describe a system for task description and simulation that enables the evaluation of these complex interdependent reliabilities. Whereas our previous work has been primarily concerned with the development of the methodology, this paper presents experimental results from applying that methodology to answer important design questions in the multirobot domain.

The only known work preceding ours in the area of predicting robot team reliability is [6]. That paper's methods are similar to ours in that they are based in the reliability literature, but that paper 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. That paper also makes comparisons only in terms of the amount of work that can be completed by different robot teams, while our methodology is built around the concept of mission tasks, which will allow us to more easily integrate our work with existing mission planning systems, most of which consider a mission as a collection of tasks.

II. Types of Failure Being Addressed

Many factors can cause the failure of a robotic mission. The laboratory robots with which most researchers are familiar usually fail due to design, manufacturing and usage errors. The hardware breaks down due to being poorly designed or constructed; the software has bugs that are revealed only under the stress of a demo; and both hardware and software fail because the robots are used in situations beyond the intentions of their designers.

While these types of failures are significant, and in fact are the dominating failure modes for most robots today [7,8,9], we contend that these failure modes are not in need of modeling so much as they are in need of correction. These failures are the result of errors and can be reduced if not eliminated through process control. Methods for reducing errors in design, manufacturing, software development and operation are widely used in industry. As mobile robots become more common, these engineering and manufacturing methods will be applied to them, yielding a reduction in these types of failures.

We can see that this is possible because some of today's robots are already built with a high degree of quality control in design, construction, and operation. For instance, the planetary rovers built for NASA by the Jet Propulsion Laboratory are built to very high standards of quality and controlled by highly trained operators, resulting in a very low incidence of failures due to errors. Another example is autonomous aerial robots. Even in the university environment, aerial robots demonstrate considerably higher reliability than ground robots. This is largely because much greater care is given to their design, construction, and operation due to the more severe consequences of failure in comparison with ground robots.

When failures due to errors are largely eliminated, as with the NASA rovers, then the remaining failures are due mostly to physical properties of the materials and to the processes used. An example of such a failure is the degradation of the grease in a bearing and the subsequent failure of the bearing. There is no process control that will change the physical reality that grease breaks down and ungreased bearings fail. Instead, the system must be designed taking into account the possibility of bearing failure.

It is this latter type of failure with which we are concerned in this paper. The reliability engineering literature provides well-established probabilistic models for this type of failure. It is possible that some of the other types of failure mentioned above can be modeled probabilistically and incorporated into these predictions. For instance, predictive models for generation of software errors have been proposed in the literature (e.g., [10,11]). Incorporation of such models would allow us to provide a more complete picture of robot failure. However, these models have been in existence for a much shorter time than hardware reliability models and have been applied in very few cases, so their ability to predict software failures is unproven. In addition, the input data required for these models are often not available in the early stages of a project, and it is this early design phase which our work targets.

III. Example Mission Scenario

A. Mission and Tasks

In these experiments we examine an example planetary exploration mission. In this mission a team of robots is tasked to install a solar panel array for a measurement and observation outpost. The mission consists of carrying the solar panels from the landing site to the outpost and then assembling them. The size of the solar panels is such that two robots 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. Robots and Components

The robots 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 robots are divided into the subsystems listed in Table 1.

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

TABLE 1
Robot Subsystems and Reliabilities

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

Transit and *Assemble*.

The probability of a subsystem's failing during a task is found using standard reliability engineering methods assuming a constant hazard rate. Two inputs determine the module failure probability: the module's failure rate, often given by mean time to failure (MTTF), and the length of time for which the module is operated during the task.

The failure rates for the robot subsystems were calculated from the failure rates of the major components in each subassembly and are listed in Table 1. The component reliability data used to derive these subsystem reliabilities were provided by the Jet Propulsion Laboratory and are representative of components used in NASA's planetary robots. An example component breakdown for the power module is shown in Table 2. Additional details on the calculation of subsystem failure and the combining of component reliabilities can be found in [5].

In addition to the failure rate, we must know the usage of each subsystem for each subtask. These usage times, shown in Table 3, were assigned using reasonable assumptions about the relative durations of different tasks and the relative usage of different modules.

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

$$P = e^{\left(\frac{-t}{MTTF}\right)} \tag{1}$$

where

t = the amount of time that the subsystem is used during the task; and

$MTTF$ = the mean time to failure for the subsystem.

Using Eq. (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.

C. Robot Teams

The baseline robot team consists of a pair of robots that are constructed to very high levels of robustness. These robots are composed of highly reliable components, are designed with operating limits well beyond the expected operating conditions, and incorporate redundancy and self-diagnostic capabilities. In other words, they are designed in the way that NASA currently designs robots. We use the MTTF values listed in Table 1 for this robot team, since the component failure rates used to derive these values are representative of actual NASA robots.

Against this baseline configuration, we examine the effects of varying both the number of robots on the team and the reliability of the components used. Among other things, we wish to compare the reliability of a larger team of less-reliable robots against the baseline team.

TABLE 2
COMPONENTS COMPRISING POWER SUBSYSTEM

Component	Quantity
Battery	2
Battery control board	2
Mission clock	1
Power distribution unit	1
Power control unit	1
Shunt limiter	1
Electrical heater	2
Radioisotope heater	2
Thermal switch	2

TABLE 3
SUBSYSTEM USAGE BY TASK

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

IV. 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–transition diagram as in Fig. 1. This particular diagram shows a team consisting of four robots that is tasked to install P panels.

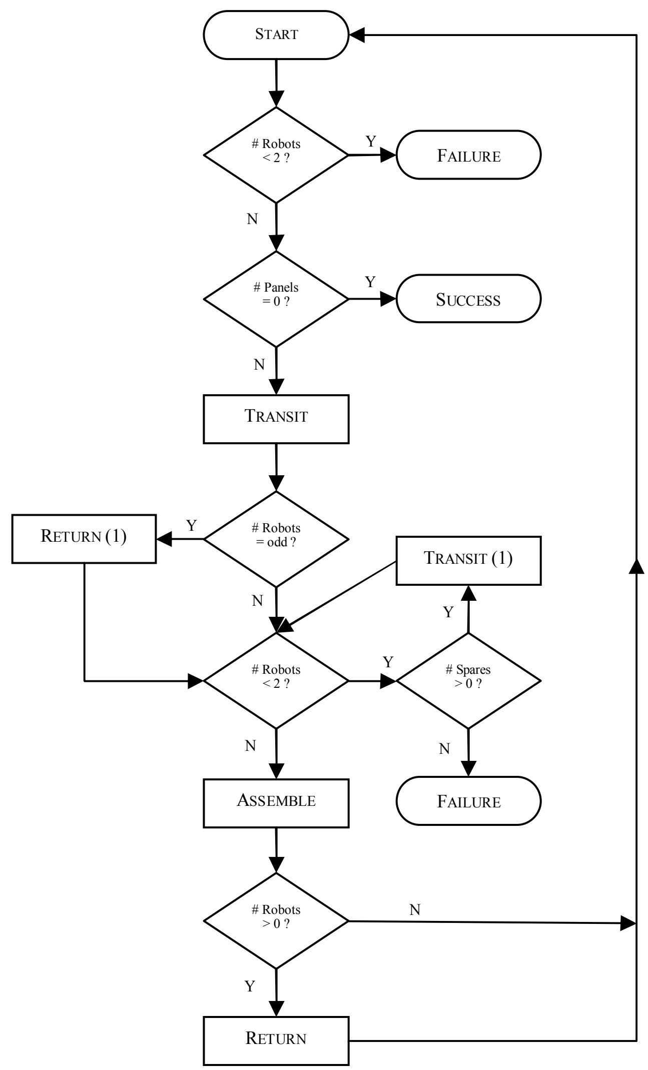

Fig. 1 State transition diagram for two-robot team.

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%

The state machine represented by the state–transition diagram is implemented in software. At each task node the state of the robot (dead or alive) is evaluated by choosing a random value between 0 and 1 for each subsystem and comparing that value with the probability of survival for that subsystem for that task. The branch in the diagram corresponding to the resulting team state is followed, and the process continues until the system reaches either *Success* or *Failure*.

As an example, after the assemble task, we would "roll the dice" for each module for each robot and compare the values with the probabilities in Table 4. If at least one of the robots survived this task, then the main branch of the diagram in Fig. 1 is followed; i.e., the *Return* task is performed. Otherwise, the diagram branches back to *Start*, since there are no robots to *Return*.

The simulation is repeated many times, with each *Success* result being assigned a score of 1 and each *Failure* result being assigned a score of 0. The average score of all trials then gives the overall probability of mission completion. The results of the simulations were verified by hand calculation for a few simple cases.

V. RESULTS

For the example mission scenario described above, once the tasks, the task durations, and the baseline module reliabilities are fixed, then the input variables for the model are:

- the number of robots on the team,
- the reliability of the components used, and
- the mission duration (number of panels to be installed).

Two of the questions that a mission designer might want to ask when choosing robots for this mission are:

"For a given mission duration and component reliability, what is the fewest number of robots that will meet a certain probability of mission completion?" and

"If additional robots are added beyond the minimum number, can we use lower reliability components, and if so how much lower?"

A. Minimum Number of Robots Required

Our initial comparison is of teams using different numbers of identical robots. Fig. 2 shows the simulation results for teams of two to six robots over a range of mission durations.

Fig. 2 shows, for example, that for a mission specifying that 30 panels are to be installed with a PoMC of at least 95% the team must have at least four robots. This figure also shows that there is a diminishing return in terms of mission reliability as more robots are added.

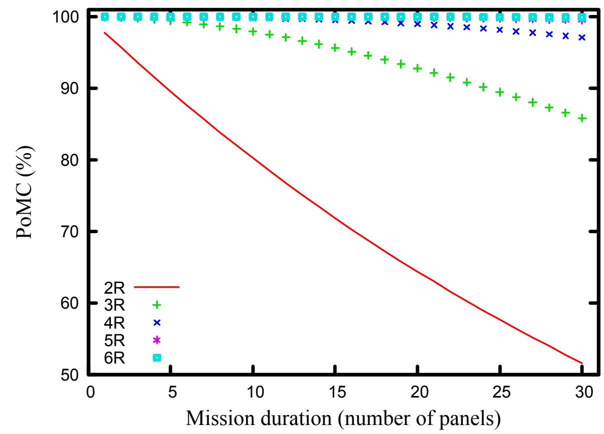

Fig. 2 Different numbers of robots with same component reliabilities.

B. Minimum MTTF with Excess Robots

If additional robots are added beyond the minimum required, it should be possible to use less-reliable components in those robots and still achieve a required mission reliability. Fig. 3 shows the simulation results for teams of four robots with component reliabilities ranging from 10% to 50% of the baseline amounts from Table 1.

When varying the reliability of the components, we apply a constant multiplier to all of the MTTF values in Table 1. For instance, when we refer to a team with 10% of the MTTF of the baseline team, we are multiplying all the values in Table 1 by 10%.

Fig. 3 shows that for very short missions a team of four robots with only 10% of the reliability of the baseline team can provide a higher PoMC. As the length of the mission increases, the reliability required for the four-robot team to equal the performance of the baseline team increases, but even for fairly long missions, the four-robot team can still outperform the baseline team even with a much lower MTTF.

To answer the question posed above—"How much lower can the reliability of the components for the four-robot team be?"—we need to look at the intersections of the four-robot curves with the two-robot curve in Fig. 3. These points give

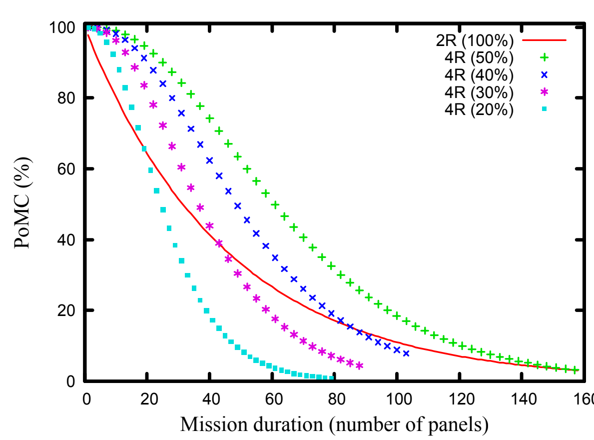

Fig. 3 Different component reliabilities.

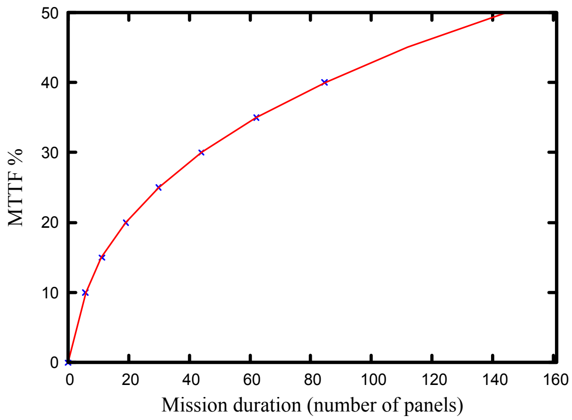

Fig. 4 MTTF % for which four-robot and two-robot teams have equivalent PoMC.

the MTTF % for which the two teams provide the same PoMC. These intersection points are replotted in Fig 4 on a graph of MTTF % versus mission length. We have also fitted a curve to these points, allowing this equalizing MTTF % to be found for intermediate points without running additional simulations.

Fig. 4 shows, for instance, that if we are designing a mission to install 20 panels, then the team of four robots will need an MTTF approximately 20% of the baseline in order to provide the same PoMC as the baseline team.

Looking back at Fig. 3, we observe that at the points of intersection the slope of the four-robot team is always steeper than that of the two-robot team. This means that the performance of the four-robot team will be more susceptible to errors in the estimates of mission parameters.

As an example, consider the 20-panel mission, for which the PoMC at the intersection point for the four-robot team with 20% MTTF is about 63%. If during the mission the assembly operation ends up taking 25% longer than anticipated, then by running new simulations with this change we find that the PoMC for the baseline team drops to 61% while the PoMC for the four-robot 20%-MTTF team drops to 56%.

A mission designer would need to take these slopes into account when selecting team configurations and components. If there is a large amount of uncertainty in the input parameters, it may be necessary to overdesign the four-robot team to a greater extent than would be necessary for the two-robot team. This may change the preferred team type in some situations.

C. *Time Required*

When choosing among robot team configurations it is necessary to consider other performance metrics besides PoMC. For instance, it may sometimes be preferable to choose a team configuration that provides a lower time to complete the mission, even if that configuration has a lower PoMC.

For the mission analyzed here, larger teams will complete the mission more quickly, since they can perform the work in parallel (assembling more than one panel at a time). Fig. 5 shows the average number of hours per completed panel for the baseline and four-robot 50%-MTTF teams. The hours-per-panel for the baseline team is simply the total time required for the *Transit*, *Assemble*, and *Return* tasks. The hours-per-panel for the four-robot team starts at half this value and climbs upward with increasing mission duration but is still significantly lower than the baseline team even for 150 panels.

D. *Cost*

Another important factor in choosing a team configuration is cost. Lower-reliability components should cost less than those with higher reliability. For a given mission, we would like to be able to determine which team configuration will provide the required reliability at the lowest cost.

In choosing components for a mission, the designer would make choices among a small number of alternate 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. Ref. [12] provides a general model for this relationship, which is given as

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

where

R_i = a reliability of interest between R_{min} and R_{max};

c = the relative cost of R_i compared to R_{min}; and

f = the feasibility of reliability improvement (between 0 and 1).

Using (2) with a feasibility of 0.5, we find that for the sample mission a cost reduction of 50% can be accomplished by choosing components with MTTF that is 40% of the baseline values. Therefore, a team of four robots with 40% MTTF would cost approximately the same as the baseline two-robot team.

Looking at Fig. 3, we see that the 4R, 40% team has a higher mission reliability than the 2R team for missions shorter than 85 panels, so the 4R team would be the more cost-effective solution for missions shorter than 85 panels.

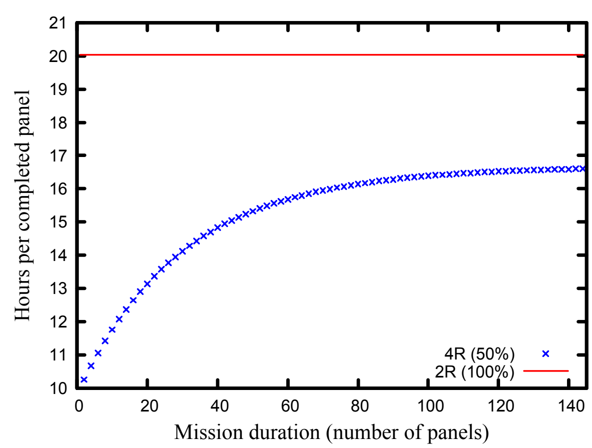

Fig. 5 Average time required per completed panel.

E. Partial Success

An additional consideration is that for many missions, including the current one, a binary representation of mission success may not be completely appropriate. For the current mission, installing a solar panel array is not an all-or-nothing venture. If only some of the panels are installed, the array will likely still be able to provide useful energy. Fig. 6 compares teams in terms of the average percentage of the assigned panels which are successfully installed. This figure shows that the four-robot teams have an even greater advantage over the two-robot team when partial mission completion is acceptable.

VI. Summary and Future Work

We have shown in this paper how reliability can guide the design of multirobot missions. Our results in this paper are significant because they provide the first quantitative support for the argument that larger teams of less-reliable robots can provide superior mission reliability compared to smaller teams of more-reliable robots, at least for some missions.

For the simple mission analyzed here, our results show that a team of four robots can provide a higher probability of mission completion than a team of two robots, even when the team of four is made of components of much lower reliability. For short missions, the four-robot team can use components with an order of magnitude lower reliability and still provide higher mission reliability. Even for fairly long missions, a four-robot team using robots with 40% of the reliability of those in the two-robot team still provides better performance. Using a parametric estimate of the cost–reliability relationship taken from [11], we have shown that the four-robot team can deliver higher mission reliability at lower cost than the two-robot team.

In future work, we plan to integrate these reliability estimation methods with mission planning software, in order to provide tools that a mission designer can use to make informed tradeoffs between mission reliability and other factors such as cost.

In addition, we intend to improve the reliability model by removing some of the simplifying assumptions currently used. For instance, we would like to allow for consideration of partial failures of robots rather than simply using the current binary dead-or-alive model. In a complex mission scenario with heterogeneous robots performing heterogeneous tasks, the failure of a robot subsystem may not render that robot useless but may instead result in re-assigning that robot to different tasks.

A number of new questions are raised when we consider how robot failure affects task allocation and re-allocation, such as "Is it better to re-assign a partially-failed robot to a new task, or to abandon it?" and "How should the initial assignment of tasks be made such that individual robot failures will have the lowest impact on the overall mission?"

Ultimately, we would like to apply these tools to a large variety of missions in order to determine if generalizations can be made about the suitability of certain types of robot teams for certain missions. We wonder, for instance, if there are classes of missions for which it is always better to use a single (or a few) highly-robust robots, and other classes of missions for which it is always better to use larger numbers of less-robust robots.

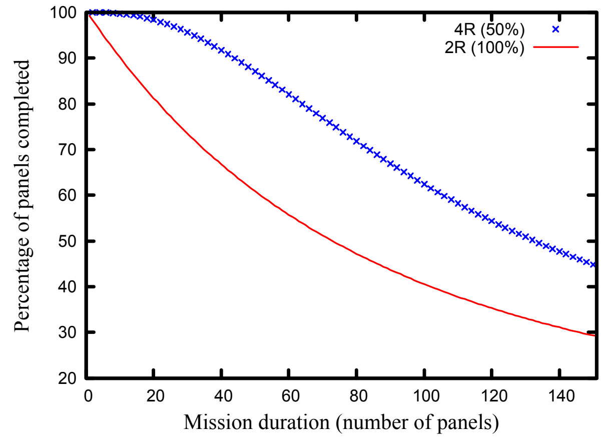

Fig. 6 Panel completion percentage.

References

[1] R. Brooks and A. Flynn, "Fast, cheap and out of control: A robot invasion of the solar system," *J. British Interplanetary Society*, October 1989, pp. 478–485.

[2] L. Chaimowicz, V. Kumar, and M. Campos, "A framework for coordinating multiple robots in cooperative manipulation tasks," *Proc. SPIE 4571–Sensor Fusion and Decentralized Control IV*, Boston, November 2001, pp. 331–336.

[3] L. Parker, C. Touzet, and D. Jung, "Learning and adaptation in multi-robot teams," *Proc. 18th Symp. on Energy Engineering Sciences*, 2000, pp. 177–185.

[4] M. Dias and A. Stentz, "A market approach to multirobot coordination," Technical Report CMU-RI-TR-01-26, Robotics Institute, Carnegie Mellon University, 2001.

[5] S. Stancliff, J. Dolan, and A. Trebi-Ollennu, "Mission reliability estimation for repairable robot teams," *Int. J. Adv. Robotic Systems,* vol. 3, no. 2, June 2006, pp. 155–164.

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

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

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

[9] J. Carlson and R. Murphy, "How UGVs physically fail in the field," *IEEE Trans. Robotics*, June 2005, pp. 423–437.

[10] Rome Laboratory, "Methodology for software reliability prediction and assessment," Technical Report RL-TR-92-52, vol. 1 and 2, 1992.

[11] M. Yin, L. James, S. Keene, R. Arellano, and J. Peterson, "An adaptive software reliability prediction approach," *Proc. 23rd Annual NASA/IEEE Software Engineering Workshop,* NASA Goddard, 1998.

[12] A. Mettas, "Reliability allocation and optimization for complex systems," *Proc. IEEE Annual Reliability and Maintainability Symp.*, 2000, pp. 216–221.