PLANNING TO FAIL:
MISSION DESIGN FOR MODULAR REPAIRABLE ROBOT TEAMS

Stephen B. Stancliff[(1)], John M. Dolan[(1)], Ashitey Trebi-Ollennu[(2)]

[(1)]The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA, stancliff@cmu.edu, jmd@cs.cmu.edu
[(2)]Jet Propulsion Laboratory, Pasadena, CA 91109 USA, Ashitey.Trebi-Ollennu@jpl.nasa.gov

ABSTRACT

This paper presents a method using stochastic simulation to evaluate the reliability of robot teams consisting of modular robots. For an example planetary exploration mission we use this method to compare the performance of a repairable robot team with spare modules versus nonrepairable robot teams. Our results show that for this mission a repairable robot team can provide a higher probability of mission completion than a nonrepairable team, even when the nonrepairable robots are built using components with an order of magnitude greater reliability than the repairable robots.

1. INTRODUCTION

The NASA Exploration Systems, Human & Robotic Technology (H&RT) Formulation Plan identifies Strategic Technical Challenges which "must be surmounted to enable sustainable future human and robotic exploration" of our solar system [1]. These include robotic networks, modularity, reconfigurability, reusability, and redundancy. The plan further identifies the need for Intelligent Modular Systems enabling safe, affordable, effective, multifunctional robotic technologies for sustainable human and robotic exploration to meet the U.S. National Vision for Space Exploration.

Modularity, reconfigurability, reusability, and redundancy add new complexity to the mission design process for robotic exploration. Decisions must be made about how to divide tasks among robots, how many robots to use, and how to configure individual robots in order to accomplish individual tasks and overall mission goals.

A significant factor in making these decisions is the impact of robot failures on mission completion. The literature (e.g., [2]) indicates that terrestrial field robots have poor reliability, with robots being unavailable approximately half of the time.

In contrast, the planetary rovers built by NASA have very high reliability, but this reliability is achieved at very high cost. Since planetary robots operate in a poorly modelled and unstructured environment the operating lifetime of components is uncertain. To accommodate this uncertainty most planetary robotic systems are designed and tested to performance standards far beyond the mission requirements. This approach can be cost prohibitive for robotic systems operating for long periods on remote planets.

In order to send teams of robots to Mars for extended missions cost-effectively, we must consider alternate robot and robot team design paradigms. Redundancy (of robot components and of robots within teams) and repairability are two options which have been proposed to reduce the impact of component and robot failures on mission success. A common assumption behind these proposals is that redundancy and repairability will reduce mission costs by allowing the use of lower-cost components. In order to evaluate this assumption, it is necessary to quantify the impact of reliability on mission success.

The existing mission planning literature deals with robot failure primarily in terms of recovery after failure occurs (e.g., [3], [4]). Our work differs in that we are developing methods to predict failure rather than to respond to it. Our methods can be used to augment existing mission planning systems by providing estimates of failure rates in the early stages of mission design, allowing one to "plan to fail" instead of dealing with failure reactively. A mission designer will, for instance, be able to evaluate the increased probability of mission failure when lower-cost components are used and will be able to compare the relative cost of using a larger number of low-reliability robots and spares versus using a smaller number of high-reliability robots and spares.

The only known previous work studying how cooperative repair impacts the reliability of robot team missions is [5]. That paper's methods are similar to ours in being based in the reliability engineering literature, but significantly different in assuming that repair incurs no cost in terms of time and reliability. We contend that in most cases this cost of repair is significant—the robots executing the repair must delay their assigned tasks in order to perform a repair, and the act of repair increases their own chance of failure.

Additionally, [5] considers only cannibalistic repair, where all replacement parts are scavenged from failed robots, and all spares are carried by the surviving robots. Our method has been designed to be flexible with respect to repair method. While in this paper we consider only repair where spare modules are available at a central location, we intend to compare other repair options in future work.

This paper follows our previous work in [6] and [7]. In [6] we presented a method for quantifying the reliability of robot modules and individual robots given component reliabilities. In [7] we demonstrated how module and robot reliabilities can be used to evaluate mission design alternatives for a simple mission. The method in [7] is labor-intensive and is therefore suitable for evaluating only very simple missions. In this paper we present a more automated method which is suitable for evaluating more complex missions.

In the remainder of this paper we first lay out an example mission scenario, then describe our method and assumptions, and finally present a comparison of repairable and nonrepairable team performance for the example mission.

2. MISSION SCENARIO

A team of robots is tasked to install a solar panel array for a measurement and observation outpost. The installation task consists of carrying the solar panels from the drop zone 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. Once the robots reach the outpost with the solar panels they follow a carefully sequenced and choreographed set of deployment steps.

We consider first a pair of nonrepairable robots that are constructed to very high levels of robustness using heritage-based design, i.e., heuristics extrapolated from previous missions. 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.

The nonrepairable robots can recover from certain pathological faults on their own. Other such faults require the intervention of ground operators. Catastrophic failure of components may lead to early termination of the mission. In the case of the robotic arm, an actuator failure would compromise safe deployment of the solar panels. In the case of drive or steering actuators, failure may require several times the nominal number of steps for precise placement of solar panels and may also require ground operators in the loop.

Against this baseline configuration we consider two alternate team configurations, each using robots that are designed to lesser standards of robustness. One configuration uses spare robots, i.e., more than two nonrepairable robots, and the other uses two repairable robots with spare components. These configurations have the advantage of being able to continue the mission after catastrophic failure of robot components, either by replacing the failed robot component, or by replacing the entire failed robot.

We assume that the launch capacity for the mission is fixed and is equal to the capacity required to transport the two highly reliable robots. Because high reliability usually implies large size and weight (due to redundancy and robustness of components), we hypothesize that we can carry more than two robots using the same launch capacity if the robots are designed with less redundancy and less-robust components. Therefore we can compare these robot team configurations within the context of equal launch capacity and thus equal launch cost. We expect that the development costs for the repairable robots will be as much as twice the cost for the nonrepairable robots due to the need to implement new technologies as well as the lack of heritage designs.

3. ANALYSIS

3.1 Robots and Modules

In our analysis we treat both repairable and nonrepairable robots as being constructed of multiple hardware modules, as in Fig. 1. A robot fails when one of its constituent modules fails. For nonrepairable robots, failure is terminal. For repairable robots, the failed module can be replaced by a spare module if one is available. The module replacement procedure is carried out by a robot other than the failed robot.

We hypothesize that with the same launch capacity we can carry either two of the highly-robust robots, three of the less-robust robots, or two of the less-robust robots plus four spare modules. We refer to these three team configurations as 2NR (two nonrepairable robots), 3NR (three nonrepairable robots), and 2RR+4M (two repairable robots plus four spare modules).

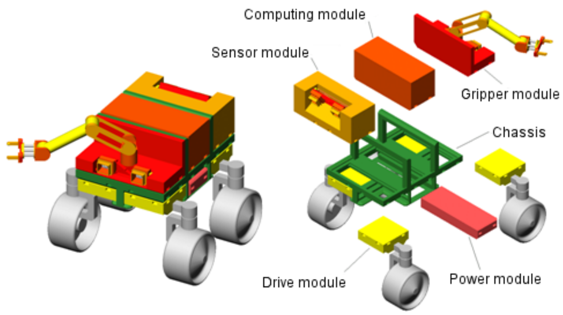

Figure 1 - Modular robot concept.

3.2 Tasks and Reliabilities

We break the task of installing a single solar panel into three subtasks:

- *Transit* to the outpost (carrying the panel)
- *Assemble* the panel
- *Return* to the drop zone

To simplify the analysis, we consider the *Transit* and *Return* subtasks to be the same. In the real world there may be some difference in reliability for these two tasks due to the extra load carried during the *Transit* subtask or due to environmental conditions (perhaps the terrain slopes towards the outpost.)

The probability of a module failing during a subtask 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 the module is operated. More details on the calculation of module failure probability are given in [6].

In this analysis 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.

3.3 Repair

Spares (either robots or modules) are stored at the drop zone. For the repairable team, when a failure occurs, the robot executing the repair must first retrieve a spare module from the drop zone and then return to the failed robot to execute the repair. This is a significant departure from the repair scenario in [7] where all of the spare modules were carried by the robots. This adds new complexity to the problem because now we must consider the location of the robots when failure occurs. An additional difference is that in that paper we were not differentiating between types of spare modules, where here we consider separately how many spares of each type are available.

Because we consider failure only at the ends of tasks, robots will always fail either at the drop zone or at the outpost. There are therefore two different repair actions to be considered for the repairable team. If a failure occurs at the outpost, the robot executing the repair must transit to the drop zone, retrieve a spare module, return to the outpost, and then execute the repair. If the failure occurs at the drop zone, then the robot only needs to pick up a spare module and execute the repair.

One simplification we make in our analysis of repairable teams is the assumption that the robot executing the repair can carry as many modules as needed for a repair. This may be unrealistic depending on the size of the robots and modules, and therefore may bias the analysis in favor of the repairable team.

Another simplification is the assumption that the cost of all repair actions is the same. In the real world this assumption may not hold. For instance, more effort or time may be required to replace one module than another. Another example is that the effort and time required for replacing two modules together would probably be less than the effort for replacing the two modules separately.

For the nonrepairable team, the "repair" action (a new robot replacing the failed one) is also different depending on the location of the failure and which task was most recently completed. With one spare robot there are two possible repair actions. If a single robot fails after the transit task, then the replacement robot will need to transit to the outpost to perform the assembly task. If a single robot fails after the assembly task, then the replacement robot will wait at the drop zone for the surviving robot to return.

3.4 Method

Our method in [7] was to enumerate all the possible paths by hand for several alternate team configurations, determine the equations representing the reliability for each configuration, and then extrapolate to a general equation for all possible configurations. This was a labor-intensive process even for the very simple mission scenario considered in that paper.

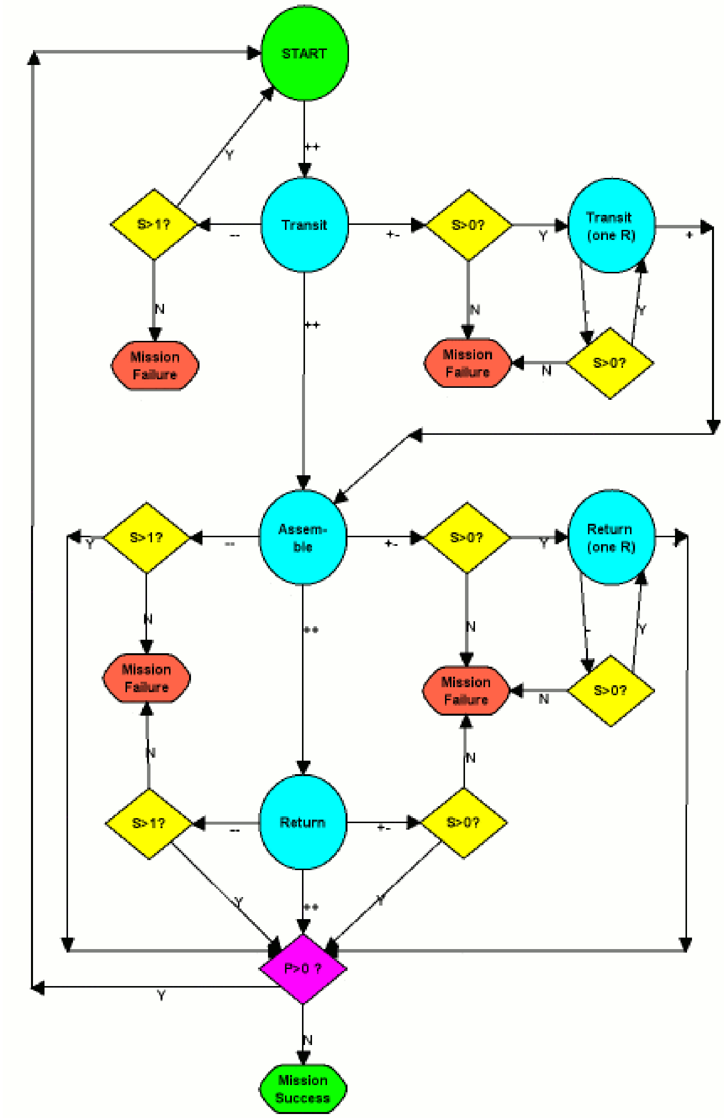

Fig. 2 - State transition diagram for nonrepairable team.

The current mission is still simple enough that an analytical solution is possible, although such an analysis would be considerably more involved than the example in [7]. However, our end goal is to provide tools for the analysis of complex missions and we want those tools to be usable by a mission designer without having to perform tedious derivations. Therefore, in this paper we have developed a new method for solving these problems using stochastic simulation. This method requires considerably less work on the part of the mission designer. The formalism for this new method remains consistent with our prior method of analysis.

In order to use stochastic simulation we have represented the mission using state transition diagrams. Fig. 2 shows the state transition diagram for a nonrepairable team with S spare robots installing P panels. The diagram for repairable robot teams is similar. The transitions from task to task are governed by the states of the robots, with "+" representing "alive" and "-" representing "dead." For most of the tasks there are two robots involved, so there are four possible robot team states: "++", "+-", "-+" and "--". In the current example the system is symmetric so that "+-" and "-+" are equivalent.

We implemented the state machines represented by these diagrams in software. The state of the robot team is evaluated at each task node by choosing a random value between 0 and 1 for each module and comparing that value with the probability of survival for that module for the current task. The branch in the diagram corresponding to the resulting team state is followed, and the process continues until the system reaches either *Mission Success* or *Mission Failure*.

The simulation is repeated many times, with each *Mission Success* result being assigned a score of one and each *Mission Failure* result being assigned a score of zero. The average score of all trials then gives the overall probability of mission success.

For the results shown here we found that 200k trials were sufficient to give repeatability to the third decimal place. The computing time required is roughly linear in the number of trials and the number of panels. For instance, on a Pentium 4, 2.6GHz, running 200k trials for 5 panels takes 2.9 s and running 200k trials for 10 panels takes 5.6 s.

In order to verify that the state transition model matches our previous model and that we have implemented the model correctly in software, we hand-calculated the results for a few simple cases using the methods from [7] and found the simulation results to be in agreement with the hand calculations to the third decimal place.

4. RESULTS

4.1 Base Comparison

In the results presented here we consider the robots to be composed of three modules:

(a) computation / power,
(b) mobility, and
(c) manipulation.

The reliability of each module is given by it's mean time to failure (MTTF). For the 2NR team, we use hypothetical MTTF values of:

(a) 40000 h,
(b) 20000 h, and
(c) 10000 h.

For the 3NR and 2RR+4M teams we hypothesize using modules with 1/10 the reliability of the 2NR modules, or:

(a) 4000 h,
(b) 2000 h, and
(c) 1000 h.

Module usage for each task is given in Table 1. Note that the duration of the *Transit*, *Assemble*, and *Repair* tasks are 6, 8, and 2 hours respectively.

Table 1 - Module usage by task

Task	Module a	Module b	Module c
Transit	6 h	6 h	0 h
Assemble	8 h	4 h	8 h
Repair	2 h	1 h	2 h

For the repairable teams, the spares available are one each of modules a and b, and two of module c. Looking at the MTTF values this seems intuitively to be the best assignment of four spares across the three module types, and was verified to be so by trial and error comparison with other assignments.

Fig. 3 compares the probability of mission completion for the three robot teams using the method and assumptions given in Section 3, for different numbers of panels to be deployed.

We see from Fig. 3 that the repairable robot team provides a higher probability of completing the mission, even though it uses components with much lower MTTF. We also see that the team with three less robust robots is better than the team of two superior robots when the number of panels is less than six, but that this team's performance drops off much more rapidly than the other two teams as the number of panels increases. It makes sense that the 3NR team will have a much lower reliability than the others when the number of panels is large, since it can only recover from a single robot failure.

4.2 Effect of MTTF ratio

Our hypothesis that the component reliability for the

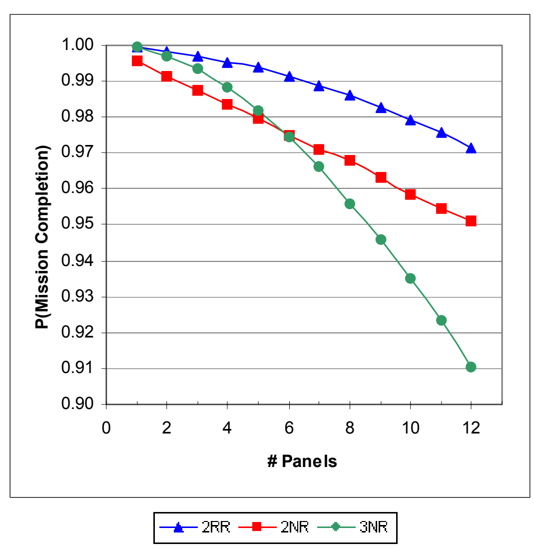

Fig. 3 - Base comparison.

less robust robots would be 1/10 that of the highly robust robots may not be accurate. In order to see how this ratio affects the repairable team performance, we reran the repairable team simulation for ratios of 1/20 and 1/5. The results are shown in Fig. 4. This figure shows that when the MTTF for the repairable team is much smaller than the 2NR team, the repairable team has a small advantage for very few panels, but the performance of the repairable team drops off rapidly. As the relative MTTF of the repairable team increases, the repairable team both has a larger initial advantage, and also its performance drops off more slowly with more panels.

4.3 Effect of Operating Conditions

We are also interested in knowing how the comparison between team configurations will change as operating conditions change. For instance, if the terrain the robots must traverse turns out to be more challenging than expected, this may reduce the effective MTTF of the mobility modules and may increase the amount of time required to complete a traverse. We discuss the effect of operating conditions on MTTF in [6].

In order to compare the teams under deteriorated operating conditions, we repeat the simulations with the same conditions as in Fig. 3 except that the MTTF for the mobility modules is reduced by half. The results are shown in Fig. 5.

We see from Fig. 5 that under these deteriorated operating conditions, the 3NR team is still clearly the worst performer when there are more than a few panels. What is different is that the 2NR team now outperforms the 2RR+4M team when there are more than eleven panels. It makes sense that the repairable team in this example mission will suffer more from worsened operating conditions, since the repairable team must complete additional transits in order to execute repairs, and the reliability of each of those extra transits is worsened. We expect that a repairable team that carried spare modules on the robots would require fewer extra transits and therefore would be less hindered by deteriorated operating conditions.

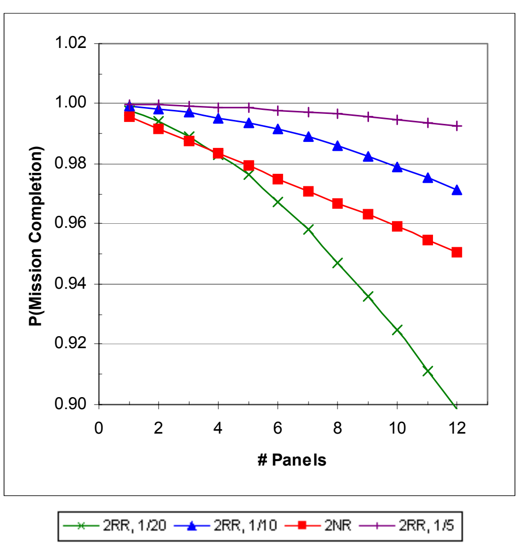

Fig. 4 - Effect of MTTF ratio.

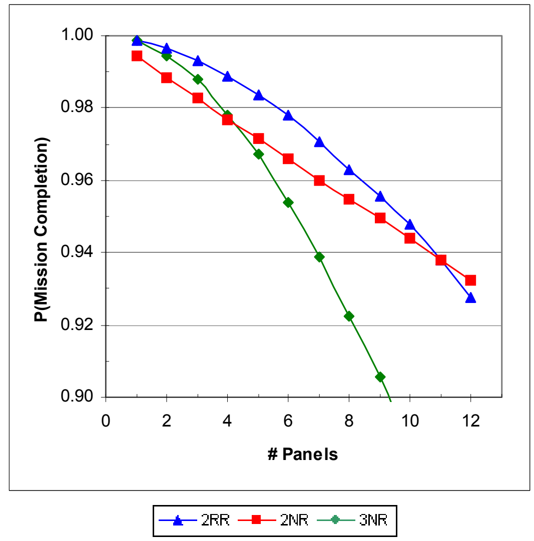

Fig. 5 - Effect of MTTF reduction.

4.4 Number of Panels Assembled

For the task of installing a solar panel array, it may not be correct to assess the mission as completely failed if the team fails to install all of the panels. The array may be incomplete, but it will likely still be useful. It is therefore interesting to examine how many panels each team installs on average, regardless of whether the entire mission is completed. This is plotted in Fig. 6 as the percent of the assigned panels which are successfully assembled, using the same conditions as Fig. 3. From this figure and Fig. 3 we see that the 2RR team not only provides a greater probability of completing the entire mission, but also provides a greater percent of panels installed when both complete and incomplete missions are considered.

4.5 Mission Delay

Finally, we want to examine the extent to which the mission is delayed by repair. For time-critical tasks the additional reliability of the repairable team may not be useful if it comes at the cost of significant delay.

We look at the time spent per panel assembled. For every installed panel there will be exactly one

Assemble subtask, so we can disregard that time in the comparison. Also, the total time spent on the *Repair* subtasks is small compared to the total time spent on *Transit* subtasks, so for a rough estimate of time spent per panel we can simply use the average number of transits required to assemble a panel. This is shown in Fig. 7 for the same input parameters as in Fig. 3.

We see from Fig. 7 that the repairable team has a fairly constant penalty of about 3% compared to the 2NR team.

5. SUMMARY

In this paper we have presented an improved method for predicting the probability of mission completion for teams of repairable and nonrepairable robots. This method eliminates the tedious derivation required by our previous method and is therefore more suitable for evaluating complex missions.

For the example mission studied here, our results show that a team of two repairable robots with four spare modules can provide a higher probability of mission completion than a team of two nonrepairable robots, even when the nonrepairable robots are constructed using modules which have an order of magnitude greater MTTF than the modules used by the repairable robots. The tradeoff is that the repairable team will take slightly longer on average to complete the mission.

This result is significant because an order of magnitude difference in reliability represents a significant difference in development costs. The lower cost of the components used for the repairable robots will at least partially offset the additional cost required to implement repair capability.

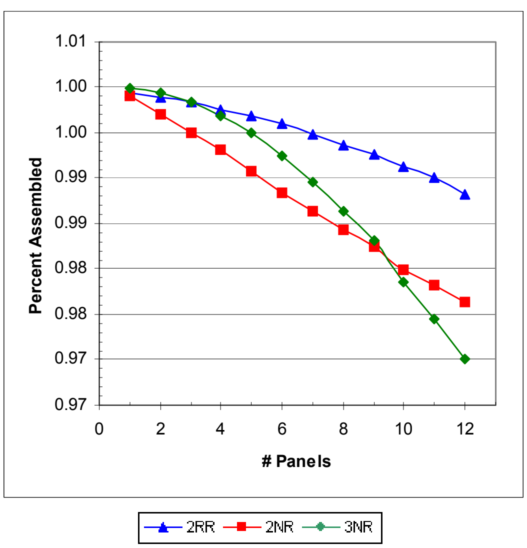

Fig. 6 - Percent of panels assembled.

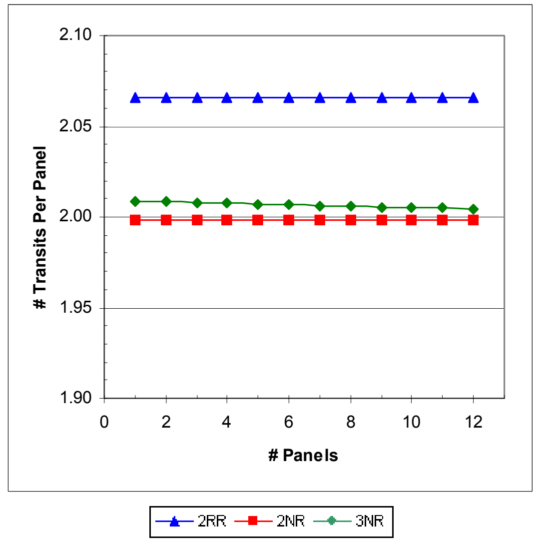

Fig. 7 - Average # transits per assembled panel.

One area for future work is to compare a variety of repair methods. For instance, in this paper all spare modules were kept at a central depot. An alternative arrangement would be to carry some spare modules with the robots, in order to reduce the amount of work required to execute a repair.

We would also like to examine a variety of mission scenarios in future work. By examining multiple scenarios and multiple team configurations, we may be able to discover patterns which will allow us to make general statements about which team configurations are best for certain types of missions.

Another area for future work is to incorporate the effects of operating conditions into the mission model, by parameterizing MTTF as described in [6].

Our ultimate goal is to incorporate these methods into mission design software which will allow a mission designer to input mission specifications and alternatives, along with real reliability and cost data, and which will then automatically generate comparisons of different robot and team configurations for that mission. To do this we will need to further automate the process, including for instance the generation of the state transition model from a mission description.

6. REFERENCES

1. NASA, "Human & Robotic Technology (H&RT) Project Formulation Plan," Version 5.1, 29 July 2004 (Amended 13 September 2004) http://exploration.nasa.gov/documents/nova_29july04b.pdf

2. Carlson, J., and Murphy, R., 2003. "Reliability Analysis of Mobile Robots," in *Proc. 2003 IEEE Int'l Conf. Robotics and Automation* (ICRA 2003), pp. 274–281.

3. Howard, A., Parker, L. and Sukhatme, G., "The SDR Experience: Experiments with a Large-scale Heterogenous Mobile Robot Team," in *Proc. 9th Int'l Symp. Experimental Robotics* (ISER'04), 2004.

4. Brummit, B. and Stentz, A., "GRAMMPS: A Generalized Mission Planner for Multiple Mobile Robots," in *Proc. 1998 IEEE Int'l Conf. Robotics and Automation* (ICRA'98), pp. 1564–1571.

5. Bererton, C. and Khosla, P., "An Analysis of Cooperative Repair Capabilities in a Team of Robots," in *Proc. 2002 IEEE Int'l Conf. Robotics and Automation* (ICRA'02), pp. 476–486.

6. Stancliff, S., Dolan, J., and Trebi-Ollennu, A., "Towards a Predictive Model of Mobile Robot Reliability," to be published as tech. report, The Robotics Institute, Carnegie Mellon University, 2005.

7. Stancliff, S., Dolan, J., and Trebi-Ollennu, A., "Mission Reliability Estimation for Repairable Robot Teams," to be published in *Proc. 1st Int'l Workshop Multi-Agent Robotic Systems* (MARS'05), September, 2005.