

An Analysis of Cooperative Repair Capabilities in a Team of Robots

Curt Bererton, Pradeep Khosla
curt@ri.cmu.edu, pkk@ece.cmu.edu
Robotics Institute
Carnegie Mellon University
5000 Forbes Ave. Pittsburgh, PA
15213

Abstract— To date, very little work has investigated the benefits of repairable robots. Robots that can repair themselves and other robots in their team are intuitively a superior design. Intuition, however, is not an acceptable basis for spending millions of dollars in development. In this work, we quantify the gain in productivity of a team of repairable robots compared to a team without repair capabilities. We create a model using an extension of standard reliability theory. It allows the definition of a metric which is used to compare the two teams. The analysis yields insight into the design of repairable robot teams under a certain set of assumptions. The model also demonstrates scenarios where repair capabilities are not likely to be beneficial.

I. INTRODUCTION

Literature on multiple robots cooperating together is in its infancy. We have just begun to be able to control a single robot well for certain tasks, let alone coordinate several. As researchers, we wish to develop tools and ideas to increase the set of tasks team of robots can perform and how well they perform those tasks. We propose that robots capable of self-repair and cooperative repair will greatly improve the performance of robots in many scenarios or tasks. To verify this theory, we have begun to study the effects of repair capabilities in robots.

Intuition indicates that robots capable of repairing each other should be superior to those that cannot. However, little is known about how this repair capability would affect the performance of a team of robots. Will these robots be far superior to other teams? Or will the cost and design effort make them impractical? In this paper, we attempt to address the following questions: *Is a team of repairable robots in fact superior to a similar team of non-repairable robots? If the repairable team is superior, by how much?*

In order to address these questions, we perform a theoretical analysis using tools from reliability theory [1], [2], [3]. Reliability theory is the study of system design with the goal of reducing the number of failures and increasing system lifetime. This theory

can be expanded to consider a system made of several robots as opposed to a system made only of components. To make the analysis tractable, several assumptions are made in generating the model. This model is used to infer several features of repairable robot teams. The general form of performance improvement is determined. The model also demonstrates cases in which the addition of repair capabilities have a negligible effect on the team's performance. The analysis also reveals some key features in the design of repairable robots which should be heeded by robot designers.

First, we review a portion of the related work followed by an introduction to the problem domain. It is shown how reliability theory can be expanded to calculate a metric with which we can compare the two teams. We then introduce our model for a team of repairable robots. The analysis then highlights the features of repairable robot teams useful to a designer.

II. RELATED WORK

There have been several works in which performance improvement due to robot capabilities is studied. In [4] Toshio Fukuda et al. study the effect of reconfiguration. They determine the performance improvements which can be obtained for a simple example task. Repair in the mobile robot domain is mostly studied with regards to reconfigurable robots [5], [6]. Reconfigurable robots have an inherent ability to repair themselves by expelling failed modules. These repair capabilities are not the typical focus of reconfigurable robot research which concentrates on the use of modularity to increase the functionality of a single robot. Two exceptions are the discussions of reconfigurable robot repair capabilities by Tomita and Yim in [7], [8]. The type of repair in self-reconfigurable systems is self-repair. There is a fundamental difference between self-repair and the cooperative repair discussed here. The difference is between carrying the necessary redundancy on the robot to repair itself or having spare parts available from other teammates. We argue that

redundancy on a single robot is in most cases a waste of resources. Redundancy used to decrease the failure rate of a given component might instead be used to send a second robot. In this work we determine the improvements possible with cooperative repair. To our knowledge, an analysis of cooperative repair has not yet been performed.

III. PROBLEM SPECIFICATION

The type of repair that we have addressed in this paper and in previous papers [9], [10] is called cooperative repair or cannibalistic repair. Robots consist of modular components which are put together to form a functional robot. In this type of repair, the functional components required to perform a repair are scavenged from robot teammates after they have failed. These can then be used to fix future failures in other teammates. We assume that there is no external entity (such as a human) to assist in this process; all repairs are to be performed by functioning robots in the team. This type of repair can be performed by the robots shown in Figure 1. These two robots, developed in [9], are capable of autonomous component removal and replacement.




Fig. 1. Robots Capable of Cooperative Repair

In order to compare repairable and non-repairable robot teams, we first discuss the choice of metric. Reliability engineering provides us with useful metrics which have long been used in manufacturing [11]. One of the most widely used metrics is the Mean Time To Failure (MTTF). The MTTF gives the best estimate of how long a particular device will function before failure. We extend this notion from a single device to the Aggregate Mean Time to Failure (AMTTF) of a team of robots. We propose that this metric should be used to compare two teams of robots.

The AMTTF is simply the sum of all the operating times of individual robots in a group or team. The AMTTF of a single robot is simply the MTTF of that

robot. If N identical robots are operating together, then the AMTTF is N times the MTTF of a single robot (assuming independence). If the robots cannot be repaired, then the AMTTF is exactly this value. In a team of repairable robots, this is not the case. By repairing a given robot after it has failed, we are essentially adding more time in which a particular robot is operating. This directly increases the AMTTF by the amount of time the newly repaired robot operates.

If there are two teams of robots whose task is to move dirt from location A to B, we would like to know which team will move the most dirt before every team member has failed. Assume each robot moves a constant amount of dirt per unit time. The AMTTF then tells us which team will move more dirt and how much more dirt it will move. If each robot performs a constant amount of work per unit time (W), then the expected total work performed by the team is: $Total\ Work = W \times AMTTF$. This brief example demonstrates why the AMTTF is a useful metric with which to compare two teams of robots.

IV. DETERMINING THE AMTTF OF THE NON-REPAIRABLE TEAM

Reliability engineering provides methods and tools to predict and evaluate the reliability of a system. One metric of reliability is the Mean Time To Failure (MTTF). This value is a single number summary of the statistical distribution representing the reliability of the system. It does not give comprehensive information about a system, but it does provide a useful metric with which two systems can be compared. Standard reliability literature [1], [2] allows us to determine the MTTF of a single robot. We then extend this notion for a team of non-repairable robots.

We assume that each robot is made of discrete components. All components of the robot must be functional in order for the robot to perform its task. The failure rates (denoted by λ) of these components are also assumed to be known. A failure rate is the number of failures per unit time of a component. Though the failure rates of most components are known to follow a bathtub type curve, we assume that the failure rates are constant. This is a common assumption in reliability literature. It is equivalent to the statement that the component has been tested past its infancy period and is not in the wearout period of its lifetime. The MTTF of a single robot with M components is given by equation 1:

$$MTTF = \frac{1}{\sum_{i=1}^M \lambda_i} \quad (1)$$

The above equation is easily extended to compute the AMTTF for a team of N non-repairable robots with

M components. As mentioned earlier, when any component of the robot fails, the robot is considered to be permanently dead and incapable of performing work. Using the definition of the AMTTF above we can then calculate the AMTTF for this team of N non-repairable robots as:

$$AMTTF_{(team)} = \frac{N}{\sum_{i=1}^M \lambda_i} \quad (2)$$

For some portions of the analysis, we force all of the component failure rates to be equal. In this case the AMTTF of the team is:

$$AMTTF_{(team)} = \frac{N}{M \times \lambda} \quad (3)$$

V. MODELLING A TEAM OF ROBOTS WITH REPAIR CAPABILITIES

A. Model Description

In order to determine the AMTTF for a team of repairable robots, we use a continuous time, stationary Markov model of the team. This is similar to Markov models used to analyze the mean time to failure of standard devices. The AMTTF can be calculated as follows. The expected waiting time in every state is calculated using the fact that the probability for remaining in the state is a negative exponential in time [1]. The number of functional robots in every state is known. The approximate intuition is that we multiply the number of robots in a state by the expected waiting time in that state. We then sum these values over all states to find the AMTTF.




Fig. 2. Markov Model for One Robot

To show how this model is equivalent to the AMTTF calculations for non-repairable robots, we present a simple example. Figure 2 shows a Markov model for a single robot. The robot is made out of one component with failure rate λ_1 . The failure rate is used as the transition rate out of the start state. Because this is a stationary, continuous time Markov process, we know that the probability of remaining in the functional state at time t is $e^{-\lambda_1 t}$. Because there is only one transition rate out of the start state, λ_1 is equivalent to ρ_i in equation 7 below. The average waiting time in this state can then be calculated using equation 6 as $\frac{1}{\lambda_1}$. This is exactly equivalent to the MTTF

of a single robot with a single component calculated using equation 1.

The MTTF values are equal because, in a team made of only one robot, the repairable team should be equivalent to the non-repairable team. This is due to the assumption that no outside components are available to repair the single robot and no external entity can directly repair a failed component. For a full discussion of how stationary and continuous time Markov processes can be applied to calculate MTTF, refer to [2] and [3].

More formally, we can calculate the MTTF of any state S_i using the following equation:

$$MTTF_{S_i} = \frac{1}{\rho_i} + \sum_{j \neq i} \frac{\rho_{ij}}{\rho_i} MTTF_{S_j} \quad (4)$$

$$\rho_i = \sum_{j \neq i} \rho_{ij} \quad (5)$$

where ρ_{ij} is the transition rate from state S_i to state S_j . The summation goes over all of the neighbors of state S_i for which ρ_{ij} is non-zero. The first term is the waiting time in state S_i . The second term says that the MTTF of S_i is the sum of the MTTF of all adjacent states S_j . Each term in this sum is adjusted by the ratio of the transition rate into state j divided by the total outflow from state i . We modify this equation to generate the AMTTF equation to be:

$$AMTTF_{S_i} = \frac{NFR_{S_i}}{\rho_i} + \sum_{j \neq i} \frac{\rho_{ij}}{\rho_i} AMTTF_{S_j} \quad (6)$$

$$\rho_i = \sum_{j \neq i} \rho_{ij} \quad (7)$$

Here NFR_{S_i} stands for the number of functional robots in state S_i . For states that transition only to failed states (i.e. states with AMTTF of 0) the AMTTF is trivially calculated as $\frac{NFR_{S_i}}{\rho_i}$.

We now expand the model for a small team of three repairable robots. We assume that each robot is made of two components. Component type one has failure rate λ_1 , component type two has failure rate λ_2 . Figure 3 shows the model for this simple case. The start state of the team is in the upper right corner where the entire team (three robots) is functional. If there are two or more robots remain in the team, then a transition to the left indicates that a type one component in one of the functioning robots has failed. A transition downward indicates that a type two component in one of the functioning robots has failed. Once there is only one robot remaining, then the only transition out of this state must lead to the state where no robots remain in operation since there are no teammates to perform a repair. The disconnected states with one and zero values are added only for completeness.

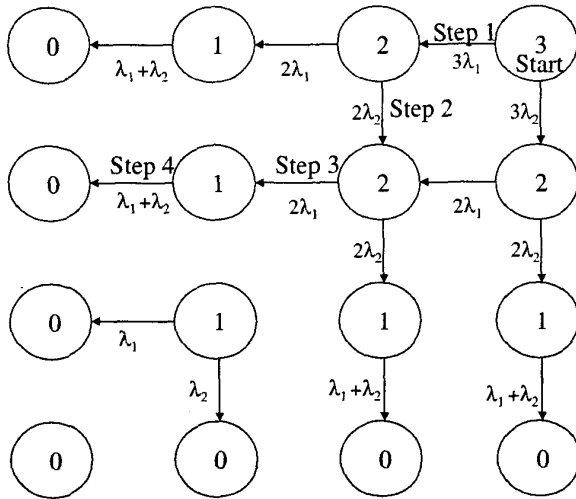


Fig. 3. Example Scenario: The Markov model is used to calculate the AMTTF of the repairable robot team.

A possible scenario for the three robot example is shown in Figure 3. In the start state, all three robots are functional. This is equivalent to the statement that there are three functional type one components and three functional type two components. In step one, a failure of a type one component occurs. At this stage there are two functional type one components and three functional type two components. This leaves only enough components for two functional robots. The next component to fail is a type two component on one of the two functioning robots. At this exact instant there are two type one components and two type two components. We make the assumption of instantaneous repair, and thus there are now two functional robots. This assumption is justified below. Another type one component fails in step three leaving only one functioning robot. Lastly, in step four, a type two component fails leaving zero robots functional. At the end of this scenario there is in fact a functional type one component and a type two component remaining. We do not allow these to form a functional robot as there is no functional robot to perform the repair.

We assume that repairs are instantaneous and that at least one functional robot is required in order for a repair transition to occur. Further, this model also assumes that no components other than those currently in use by the robots are present. No external entity is available to repair failed components nor combine functional components into functional robots. As with the non-repairable robots, we assume that individual component failures are independent and any failed component on the robot causes the entire robot to fail (though it may be repaired).

In order to include more robots we simply add a row and column of states for every additional robot. In Figure 3 for example, a row is added to the top and a column is added to the right. Thus, for four robots consisting of two components each, the number of states is 25. To account for more components, an extra dimension is added for every component. For three robots and three components, the Markov model is a cube containing 64 states. Unfortunately, this leads to a number of calculations that is exponential in the number of components. More exactly, the number of states which must be calculated given N robots made with C components each is:

$$\text{Number Of Calculations} = (N + 1)^C \quad (8)$$

The values of all one-robot states are determined first, which allows us to determine the values of all two robot states and then all three robot states, etc. Determining the value of a given start state requires iteratively calculating the values of all states which have fewer functioning robots.

VI. ASSUMPTIONS

This analysis makes several assumptions which benefit the repairable robot team. Firstly, it is not the case that repair is instantaneous. To rectify this, a delay state could be added to any state transition indicating a repair. In previous work [9], we found that repair times were small compared to overall working time. This is the justification for omitting the delay state.

Components not being used by a robot will still degrade. In the above model, we assumed that components not in use are in perfect stasis. To rectify this difficulty, the probability weight indicating a failure could be modified to include a factor for components which are not actively being used.

During a repair task, a robot may be particularly susceptible to failure. The model could be modified to include extra transitions which could include this susceptibility.

The above algorithm also neglects the fact that the robots doing the repair are not doing useful work. A component may also fail in such a way that the robot is unable to repair another robot. Losing the ability to repair others may not hamper the robot's ability to perform a task. This is certainly true if a component is only needed to perform a repair on another robot. One may also have components which only partially fail, but continue to function in a degraded capacity.

In some parts of the analysis, we assume that all components of a given robot have equal failure rates. This assumption is optimal for the repairable robot team, though it may be justified if the robots were

purposely designed in this manner. In a real robot, it is unlikely that the failure rates of each component would be equal. It is likely that the repairable robot components would have a slightly higher failure rate. They must include at least one extra connector and possibly extra components.

Failures between components are also assumed to be independent. However, there are many examples (a power surge or falling from a great height for instance) where two or more components may fail at one time.

Perhaps the most grievous of assumptions is that robot failure is due entirely to component failures. Many robot failures are due to robots simply getting stuck, or becoming lost, or batteries running out. These failures could be modelled as fake components which would cause an entire robot to fail. However, these faults would affect both repairable robots and non-repairable robots equally and what we wish to determine here is the difference between the two designs.

VII. RESULTS

The AMTTF of a variable size team of non-repairable robots with any number of components is of the form, $AMTTF_{(team)} = \frac{C}{\lambda}$, C is a constant, and λ is simply the summation of the individual failure rates. Surprisingly, the solution for the repairable robot team can be placed in the same form under the assumption that component failure rates are equal. This feature of the solution allows us to perform some comparisons between a repairable robot team and a non-repairable robot team. All of the figures and tables in subsection A and B rely on an equal component failure rate assumption. Robots in teams being compared are assumed to be made with the same number of components. In section VII-C we observe the effects of having non-equal failure rates.




Fig. 4. Two Component Repairable Robots vs. Two Component Non-Repairable Robots

A. Fixed Number of Components with Equal Failure Rates, Varying Team Size

In Figure 4, we see that for a fixed number of components and robots, the repairable robots have a higher AMTTF. Though in Figure 4 we show the results for robots made only of two components, the results for robots consisting of more components are of the same form. The curve in Figure 4 can be approximated by a line in which the slope of the line represents the AMTTF per robot. The repairable robot curve has a consistently higher AMTTF. The greater the number of robots in a given team, the greater the advantage of the repairable robot team over the non-repairable team.

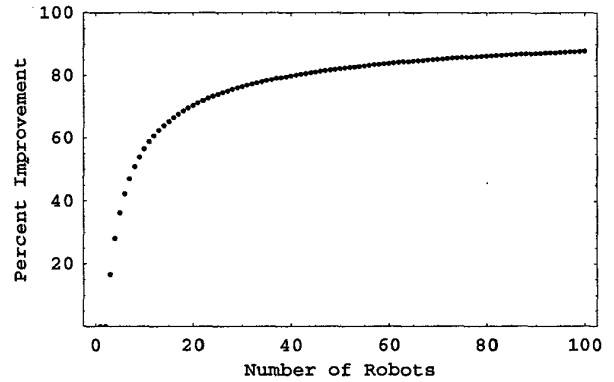


Fig. 5. Two Component Repairable Robots vs. Two Component Non-Repairable Robots

In Figure 5 below, we see the improvement of the repairable robots over the non-repairable robots as the number of robots increases. Note that for small team sizes the improvement is smaller than for larger team sizes because the first few points in the curve do not rise linearly. This is shown more clearly in Figure 6 below. The slow rise near the beginning of the repairable robot curve is not as noticeable in the six component case compared to the two component case. Our intuition is that the first robot failure yields a spare parts collection to repair future component failures with. If there is a larger probability that the same type of component will fail twice in a row (such as if there are only two components on the robot), then the repairable team will tend to fail faster.

Using a least squares fit, we determine the slopes of the curves from Figures 4, 6, as well as for similar plots of three, four, five, and seven component robots. These slopes are compared to the non-repairable robot curves generated using equation 2. From equation 2, one can see that the slope for the non-repairable robots is simply one over the number of components. The results




Fig. 6. Six Component Repairable Robots vs. Six Component Non-Repairable Robots

are shown in Table I. Note that in each case, the repairable robot slope is greater than the non-repairable robot slope.

Comparing the values of the AMTTF's of the teams, we see that the performance improvements can be substantial for team sizes greater than three or four robots. In Figure 6 we see the results for a team of fourteen robots with six components. With this team size, the repairable robots have an AMTTF five times greater than the non-repairable robots.

B. Fixed Team Size and a Variable Number of Components with Equal Failure Rates

As the number of components increases, a performance decrease comes from the fact that we are adding more components which are equally likely to fail. Table I reflects this performance decrease as a decrease in the slope of the line (amount of work done per robot) that is calculated using a least squares fit. As components are added, the repairable robot team AMTTF decreases less than the non-repairable robot team. Table I demonstrates the superiority of the repairable robot team as the number of equal failure rate components composing a given robot increases.

C. Effects of Non-Equal Failure Rates

In order to observe the effects of non-equal failure rates, we calculate the AMTTF of the start state in Figure 3. This leads to an AMTTF of:

$$\frac{18\lambda_1^2 + 49\lambda_1\lambda_2 + 18\lambda_2^2}{(\lambda_1 + \lambda_2)(3\lambda_1 + 2\lambda_2)(2\lambda_1 + 3\lambda_2)} \quad (9)$$

Similarly, using equation 2, we can determine the AMTTF of the non-repairable team as $\frac{3}{\lambda_1 + \lambda_2}$. Dividing these two equations and plotting the improvement of the repairable team over the non-repairable team versus the values of λ_1 and λ_2 yields Figure 7.

Number of Components	Work Done per Non-Repairable Robot	Work Done per Repairable Robot
2	0.5000	0.9859
3	0.3333	0.9552
4	0.2500	0.9161
5	0.2000	0.8647
6	0.1667	0.7876
7	0.1429	0.6937

TABLE I
AMOUNT OF WORK DONE PER ROBOT




Fig. 7. Improvement of Two Component, Three Robot, Repairable Team Over Non-Repairable Team

As the failure rate of either λ_1 or λ_2 goes to zero, the improvement goes to 0 percent (i.e. no improvement). The best case for repairable robots is when $\lambda_1 = \lambda_2$. This makes intuitive sense in that if one failure rate is much higher than the other, then all of these components would fail first and there would be no opportunity for repair. If the failure rates are equal, this leads to the best chance for the re-use of components to repair future failures and thus the best use of repair. The optimal design for repairable robots is therefore to make all replaceable components have equal failure rates.

VIII. ANALYSIS

From equation 9 a design trade-off is apparent. A designer may choose to spend money lowering the failure rates of components to achieve a higher AMTTF due to the higher order of λ in the denominator. The alternative is to spend money making more robots and simply sending a larger team of robots leading to a linear improvement per additional robot sent. In practice

however, it is usually very difficult to cost-effectively lower the failure rate of an existing component. They must be designed and manufactured from the ground up as a high-reliability device.

Many of the assumptions benefit the repairable robot team's AMTTF. We see in the results that small teams of repairable robots only have small performance gains over the non-repairable teams. If this is the case even with these beneficial assumptions, then we do not recommend the inclusion of repair capabilities if the proposed team size is small.

IX. CONCLUSIONS

We have presented a reliability based model of a team of repairable robots. The AMTTF is proposed as a useful metric with which to compare the repairable robot team to an equivalent non-repairable robot team. The model allows us to calculate the AMTTF of the team and perform useful comparisons of the two teams. These comparisons lead to a better design strategy. When designing a repairable robot, all components should have an equal chance of failure. This result is intuitive, however this model verifies the intuition formally.

We also determined a general case where repairable robots have negligible performance gains. Small team sizes do not provide enough incentive for repair capabilities to be added. This is the case where no external source of components is added. It may be that repair capabilities for small teams are more valuable when replacement components are readily available.

Future work consists of relaxing many of the assumptions by altering the model. Testing with the robots described in [10] will be done to determine the most common causes of failures. Once these are determined, the model can be altered to match the physical robots. We are also investigating a quantitative comparison between the redundancy techniques used in current NASA rover designs and the cooperative repair capabilities discussed here.

We believe that the investigation of repairable robots has been shown to be a worthwhile goal. The performance gains have great potential to boost productivity. There are many questions yet to be answered in terms of the cost to design and manufacture such robots. However, we feel that cooperative repair is a useful strategy to increase the performance of a team of robots and that this line of research should be pursued.

REFERENCES

- [1] Balbir S. Dhillon. *Reliability Engineering in Systems Design and Operation*. Van Nostrand Reinhold Company, 1st edition, 1983.
- [2] Alessandro Birolini. *Reliability Engineering, Theory and Practice*. Springer, 3rd edition, 1999.
- [3] Roy Billinton and Ronald Allan. *Reliability Evaluation of Engineering Systems*. Plenum Press, 1992.
- [4] Toshio Fukuda, Tomoyuki Kaga, and Kosuke Sekiyama. The meaning of functional reconfiguration under the dynamic environment and system behavior. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 1676–1683, 1996.
- [5] Keith Kotay, Daniela Rus, Masette Vona, and Craig McGray. The self-reconfiguring robotic molecule. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 1, pages 424–431, 1998.
- [6] Satoshi Murata, Haruhisa Kurokawa, and Shigeru Kokaji. Self-assembling machine. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 441–448, 1994.
- [7] Kohji Tomita, Satoshi Murata, Haruhisa Kurokawa, Eiichi Yoshida, and Shigeru Kokaji. Self-assembly and self-repair method for a distributed mechanical system. *IEEE Transactions on Robotics and Automation*, 15:1035–1045, 1999.
- [8] Mark Yim, David Duff, and Kimon Rufas. PolyBot: A modular reconfigurable robot. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 514–520, 2000.
- [9] Curt Bererton and Pradeep Khosla. Towards a team of robots with repair capabilities: Docking systems. In *International Symposium on Experimental Robotics*, 2000.
- [10] Curt Bererton and Pradeep Khosla. Towards a team of robots with reconfiguration and repair capabilities. In *IEEE International Conference on Robotics and Automation*, 2001.
- [11] Farooq Khan. Equipment reliability: A life-cycle approach. *Engineering Management Journal*, 11:127–135, 2001.