

Rolling Dispersion and Exploration for Robot Teams

Elizabeth A. Jensen

Department of Computer Science and Engineering
 University of Minnesota
 ejensen@cs.umn.edu

Abstract

After a disaster, human rescuers may have to wait for better conditions before beginning to search for survivors. A team of robots could enter long before the humans and scope out the environment to gather information that could help to prioritize tasks for the rescue operation. We have developed an algorithm to allow a small group of robots to progressively explore an unknown environment, moving as a group until full exploration is achieved. The novel concept behind this algorithm comes from the way in which the team stays together as a group, maintaining communication, in order to ensure full exploration as well as a path to the exit. We demonstrate in simulation that the algorithm works in multiple environments under varying conditions.

1 Introduction

In the aftermath of a fire or earthquake, it can sometimes take hours before the rescue team is allowed into a disaster area, due to the dangerous environment. A team of robots could be deployed into such an environment instead to explore, locate survivors, and mark structural damage. To be useful, the robots would then need to report back to the rescue team outside. It is also necessary that the robots fully explore the environment, so as not to miss any important information.

Our main contribution is a novel distributed algorithm, which guarantees full exploration, while the robots also maintain communication and exit upon completion of the task. The key feature of the algorithm is that the robots advance as a group as they explore the environment, leaving behind beacons to mark both explored areas and the path to the exit. The robots use wireless signal intensity to ensure that they stay within communication range of at least one other robot at all times. The algorithm is fully distributed, and each robot makes its own decisions on which behavior to execute depending on the situation and on the robot's current role, yet the robots operate as a team.

2 Related Work

Using a team of robots to explore an environment has become popular of late, and there are many methods available.

Gage [1992] proposed three categories of coverage—blanket, barrier and sweep coverage. In blanket coverage, the agents provide continuous coverage of every point in the environment. Barrier coverage is often used in surveillance, where a perimeter is set up around an area such that nothing can pass into or out of that area without being seen by at least one agent. Sweep coverage makes a pass over the environment and ensures that every point has been seen at least once, but the agents don't stay in any one location, instead moving progressively through the environment. Most coverage algorithms are focused on surveillance and usually entail creating a sensor network to provide either blanket or barrier coverage.

Distributed coverage algorithms are designed to scale, and can also better take advantage of the robustness inherent in the multi-robot system, making them more reliable than centralized systems, which have a single point of failure. Dirafzoon et al. [2012] provide an overview of many sensor network coverage algorithms. Many of these rely on individual robots knowing the distance and bearing to other robots around them, which requires sophisticated and expensive sensors. However, it has also been shown that a team of robots can disperse into an unknown environment using only wireless signal intensity to guide the dispersion [Ludwig and Gini, 2006]. This method allows the use of small, simple robots, which are often less expensive, making it easier to assemble a larger team. However, blanket coverage can still require a prohibitively large number of robots and some algorithms still don't achieve full coverage.

3 Rolling Dispersion Algorithm

The primary goal of the algorithm is to achieve full coverage with a small team of basic robots. This rules out blanket coverage, because the team is too small, and dead reckoning methods, because basic robots have limited sensor and computational capability. Using a distributed method allows us to take advantage of the redundancy and robustness inherent in a team of robots.

Given these goals and limitations, we established two major constraints in designing the rolling dispersion algorithm—that there are not enough robots to provide blanket coverage of the environment, and that the robots stay in communication with one another at all times. We assume that the robots can detect and avoid obstacles, communicate wirelessly, and carry and deposit beacons in the environment. We also as-

sume a disaster scenario, so the specifics of the current environment are unknown.

During the exploration of the environment, each robot will execute the algorithm and make decisions as an individual, but it will have input from the beacons and other robots. Beacons are used to mark explored areas to prevent multiple passes, the unexplored frontier, and the path to the entrance. All robots are either sentries, which stay in place and act as tethers, or explorers, which explore the frontier while staying in contact with the group. Explorers become sentries when they reach the edge of their signal range, and sentries become explorers when there are no explorers left or they are on a path that has been fully explored. The robots use the wireless signal not only for communication, but also to direct their movement. Since one goal is to maintain communication, the robots only need to know if the signal intensity is increasing or decreasing to determine which direction to move. This may not lead to the maximal dispersion, but suboptimal dispersion is acceptable since our main priority is to achieve full coverage without loss of communication.

While our algorithm uses wireless signal intensity to disperse the robots, and beacons to mark locations, the innovation in our approach lies in the manner in which the robots continue the exploration past the bounds of their initial dispersion. The robots are not allowed to move in isolation, but must always stay within communication range of the team. When there is an area to be explored that is beyond the reach of the robots nearby, then the entire team of robots will move towards the unexplored area. This approach has two main benefits. First, it means that the robots are less likely to get lost, since they will have a wireless signal to follow to get back to the entrance. Second, the robots will clear each room and corridor in a methodical manner, thus reducing the likelihood of missing an area.

4 Experimental Results

We ran our experiments in the Player/Stage [Gerkey *et al.*, 2003] simulation environment, using a cave-like environment, which requires ten robots for blanket coverage. We ran these experiments with five and eight robots, with ten runs each. It took five robots an average of 192 seconds to complete the exploration, with a standard deviation of 31 seconds. Eight robots completed the exploration in 167 seconds on average, with a standard deviation of 29 seconds. In both cases, it took an average of 13 seconds for all robots to return to the entrance after completing the exploration.

Figure 1 shows the average percentage of the environment covered over time by 5 and 8 robots, with error bars to show the standard deviation at specific time points. The eight robots start more slowly, as it takes them more time to spread apart, but then they cover the environment more quickly than the five robots. The five robots show high variability in progress between 50 and 150 seconds, which is the result of waiting on retraction when too many robots went down a short path. This did not appear to the same extent in the eight robot runs, mostly because there were enough robots that every path generally had enough robots to continue exploring while robots on another path were retracting.

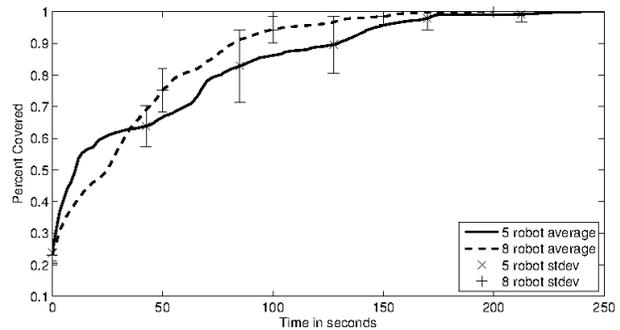


Figure 1: Average percent of the environment covered over time for five and eight robots (ten runs each). Error bars show standard deviation at that time point.

5 Conclusions and Future Work

We have developed an algorithm to disperse a small team of robots into an unknown environment to completely explore the space while staying connected at all times during the exploration. The algorithm requires fewer robots than would be needed for blanket coverage, but still provides the necessary information about the environment. The algorithm also ensures that the robots are able to exit the environment upon task completion. We have shown in simulation that the algorithm can provide complete coverage of an unknown environment.

In future work, we will test the algorithm in larger and more complex environments in simulation. We will also be running experiments with physical robots. These experiments are critical to show that the wireless signal intensity can be used to disperse the robots and maintain communication without interfering with the end goal of full coverage. In addition, we will develop formal proofs of algorithm correctness, and models for determining the most effective number of robots and beacons for a given scenario. Additional extensions to the algorithm include using the beacons to guide mobile survivors to the exit, and developing the interface and protocols for working with humans.

References

- [Dirafzoon *et al.*, 2012] A. Dirafzoon, S. Emrani, S. M. Amin Salehizadeh, and M. B. Menhaj. Coverage control in unknown environments using neural networks. *Artificial Intelligence Review*, pages 237–255, 2012.
- [Gage, 1992] D. W. Gage. Command control for many-robot systems. In *19th Annual AUVS Technical Symposium*, pages 22–24, 1992.
- [Gerkey *et al.*, 2003] B. Gerkey, R. T. Vaughan, and A. Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proc. 11th Int’l Conf. on Advanced Robotics*, pages 317–323, 2003.
- [Ludwig and Gini, 2006] L. Ludwig and M. Gini. Robotic swarm dispersion using wireless intensity signals. In *Distributed Autonomous Robotic Systems 7*, pages 135–144, 2006.