Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Multi-Robot Exploration with Communication Restrictions

Elizabeth A. Jensen
University of Minnesota
Minneapolis, MN, USA
ejensen@cs.umn.edu

Abstract

After a disaster, instability in the environment may delay search and rescue efforts until it is safe enough for human rescuers to enter the environment. Such delays can be significant, but it is still possible to gather information about the environment in the interim, by sending in a team of robots to scout the area and locate points of interest. We present several algorithms to accomplish this exploration, and provide both theoretical proofs and simulation results that show the algorithms will achieve full exploration of an unknown environment even under communication restrictions.

1 Introduction

When a human search and rescue team arrives on the scene of a disaster, it may be prevented from entering immediately due to the instability of the environment, but robots can scout the area in advance. Multi-robot systems have an inherent redundancy that increases the system's robustness [Choset, 2001], making them very useful for coverage and exploration problems. Gage [1992] proposed three types of coverage– blanket, barrier, and sweep. Most coverage algorithms are focused on surveillance, and thus aim to achieve either blanket or barrier coverage, but the number of robots required for full coverage can be prohibitively large. In contrast, sweep coverage can be done with a small team, down to a single robot, if necessary, such as when all other robots on the team have failed [Fazli *et al.*, 2010], and is more applicable when exploration is the main goal of providing coverage.

In addition to the type of coverage provided, one must consider whether a centralized or distributed algorithm would provide more benefits to the overall exploration. Distributed algorithms are designed to scale robustly, without a central point of failure, and, in an unknown environment such as ours, this features is a huge asset. Dirafzoon et al. [2012] provide an overview of many sensor network coverage algorithms which can be applied to multi-robot systems as well. Many of these rely on individual robots knowing the distance and bearing of other robots around them, which requires more sophisticated 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 simple robots, without the need to carry a heavy payload of sensors, so that the robots can run longer and explore further. Smaller, simpler robots are also less expensive, so more robots can be acquired for a task.

Our primary contribution in this work is a set of distributed algorithms for exploration using a small team of robots. The innovation in these algorithms comes from how the robots disperse into and subsequently explore the environment, even with communication restrictions and attrition. We have provided proofs that the algorithms will achieve full coverage of the environment, return all functioning robots to the entry point, and that paths to points of interest are marked for the human rescuers to follow when they are allowed entry. We show, in simulation, that the algorithms function correctly to allow the team of robots to achieve the desired results.

2 Communication-Restricted Exploration

Our primary objective is for our algorithms to achieve full exploration of an unknown environment using a team of robots. We assume that the robots used have proximity sensors to avoid collisions, some capability for communication (wi-fi, line-of-sight, chemical, etc), and the means to carry and drop off beacons. We also assume that the specifics of the environment are currently unknown, even if pre-disaster information, such as a map, is available.

The first algorithm, the Rolling Dispersion Algorithm (RDA), assumes that wi-fi communication between robots is restricted only in terms of range, which may be reduced at times by obstacles such as walls or rubble. Beacons are used to block off explored areas, or mark the path to a frontier that had to be temporarily abandoned in favor of fully exploring a different frontier. The robots initially disperse to the furthest extent of their communication range, and then call robots from other paths to move further, always keeping the team together. Further details can be found in [Jensen and Gini, 2013].

The second algorithm, the Sweep Exploration Algorithm (SEA), is based on RDA, but is intended for use in scenarios with much more restrictive communication, such as chemical signals, or line-of-sight using a camera and color LEDs. With such limited means of communication, it is critical to reduce the number and size of messages to ensure full exploration. However, this also means that only one robot can be moving at a time, or the messages get mixed up and the robots may miss a section of the environment. Therefore, instead of the robots initially dispersing in any direction as in RDA, they travel one at a time down a single path, until it is completely explored, and then retract and explore a new path. Additional details can be found in [Jensen *et al.*, 2014].

The third algorithm, the Train Exploration Algorithm (TEA), has the same communication restrictions as SEA, but rather than leap-frogging to the frontier, the robots advance and retreat as a connected group, similar to a train moving on tracks. This provides faster expansion into the frontier, because delays to wait for reinforcements are reduced to the round-trip time for the messages, not the time for the message to travel to the end of the line and a robot to make the return trip. The same is true with the retraction step, with the added benefit that each robot along the retraction path drops a single beacon, which the last robot sets to block as it passes. This also balances the number of beacons each robot has to carry, so that the last retracting robot is not responsible for dropping off all the needed beacons along the path.

We have previously presented theoretical proofs that our algorithms will, even with communication restrictions, complete the exploration without missing any point in the environment, will not end up in infinite loops (so that the robots exit when done), and can succeed in these goals even with robot and beacon failures. We have also shown that the SEA algorithm uses the minimum number of unique messages to complete the exploration [Jensen *et al.*, 2014].

3 Simulations and Results

We have conducted experiments in Player/Stage and ROS/Stage, using the same robot models and movement/sensor attributes, in order to test the viability of our algorithms. The testing environment contains several large obstacles at varying intervals and in non-uniform shapes that leave wide open areas and potential loops.

Figure 1 shows the rate of coverage for the RDA and SEA algorithms for 1, 5 and 8 robots. Both algorithms perform the same with a single robot, which shows a plateau only during retraction. RDA with 5 robots shows significant plateaus, because the robots initially disperse in all directions, and then must wait for others to leap-frog out to the frontier. The SEA algorithm starts with a higher initial coverage because a different start orientation was necessary to avoid collisions. To account for that, we compare run-times from the same starting percentage to full coverage, which shows SEA is 1.35 times faster at achieving full coverage than RDA. SEA will outperform RDA in environments where there are long paths, because SEA fully explores only one path at a time, while RDA will attempt to explore as many paths at a time as possible. However, in an environment made up single rooms off of a corridor, RDA will have the advantage, because then many rooms will be explored simultaneously, while SEA will explore only one room at a time.

4 Conclusion

We have presented here three distributed algorithms for multi-robot exploration of unknown environments. All three algorithms make use of their communication signal intensity to direct the movement of the robots, and beacons are used to mark explored areas in the environment in addition to creating a trail to the entrance and other points of interest within the environment. Each of the algorithms will allow a team of robots to fully explore the environment, so long as at least one member of the robot team remains functional.

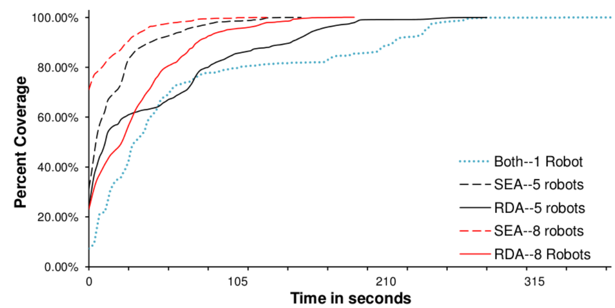

Figure 1: Average time to full exploration using 1, 5 and 8 robots with each algorithm.

In future work, we plan to test the algorithms in other types of environments, both in simulation and on physical robots, and include human interaction. An additional avenue of future work is in relaxing the restrictions on the robots' movements, in particular allowing them to move out of communication range with each other. This makes it much more difficult to prevent infinite loops and ensure full coverage, but would increase the speed of exploration.

References

[Choset, 2001] Howie Choset. Coverage for robotics–a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31:113–126, 2001.

[Dirafzoon *et al.*, 2012] A. Dirafzoon, S. Emrani, S. M. Amin Salehizadeh, and M. B. Menhaj. Coverage control in unknown environments using neural networks. *AI Review*, pages 237–255, 2012.

[Fazli *et al.*, 2010] Pooyan Fazli, Alireza Davoodi, Philippe Pasquier, and Alan K Mackworth. Complete and robust cooperative robot area coverage with limited range. In *IROS*, pages 5577–5582. IEEE, 2010.

[Gage, 1992] D. W. Gage. Command control for many-robot systems. In *19th Annual AUVS Technical Symposium*, pages 22–24, 1992.

[Jensen and Gini, 2013] E. A. Jensen and M. Gini. Rolling dispersion for robot teams. In *Proc. Int'l Joint Conference on Artificial Intelligence (IJCAI)*, pages 2473–2479, 2013.

[Jensen *et al.*, 2014] Elizabeth Jensen, Ernesto Nunes, and Maria Gini. Communication-restricted exploration for robot teams. In *AAAI Workshops*, 2014.

[Ludwig and Gini, 2006] L. Ludwig and M. Gini. Robotic swarm dispersion using wireless intensity signals. In *DARS*, pages 135–144, 2006.