Robotic Adversarial Coverage

(Doctoral Consortium)

Roi Yehoshua
Supervisors: Dr. Noa Agmon and Prof. Gal A. Kaminka
Computer Science Department
Bar Ilan University, Israel
yehoshr1@cs.biu.ac.il

ABSTRACT

Coverage is a fundamental problem in robotics, where one or more robots are required to visit each point in a target area at least once. While all previous studies of the problem concentrated on finding a solution that completes the coverage as quickly as possible, in this thesis I consider a new and more general version of the problem: *adversarial coverage*. Here, the robot operates in an environment that contains threats that might stop the robot. The objective is to cover the target area as quickly as possible, while minimizing the probability that the robot will be stopped before completing the coverage. The adversarial coverage problem has many real-world applications, from performing coverage missions in hazardous environments such as nuclear power plants or the surface of Mars, to surveillance of enemy forces in the battlefield and field demining. In my thesis I intend to formally define the adversarial coverage problem, analyze its complexity, suggest different algorithms for solving it and evaluate their effectiveness both in simulation and on real robots.

Keywords

Mobile robot coverage;adversarial coverage;adversarial modeling;motion and path planning;robotics in hazardous fields

1. INTRODUCTION

The problem of single and multi-robot coverage has been extensively discussed in the literature and many approaches to coverage path planning have been developed (see [1] for a recent exhaustive survey). The coverage problem has many real-world applications in various domains, from automatic floor cleaning and coating in facilities, such as supermarkets and train stations, to humanitarian missions such as search and rescue.

While all previous works on the coverage problem concentrated on finding a solution that completes the coverage as quickly as possible, in this thesis I consider a new version of the problem: *adversarial coverage*. Here, the robot operates in an environment that contains threats that might stop the robot. Each point in the area is associated with a probability of the robot being stopped at that point and the probabilities can vary from one point to another. The objective of the robot is to complete the given mission—to cover the *entire* target area—as quickly as possible while minimizing the probability that the robot will be stopped before completing the coverage.

The adversarial coverage problem has an intrinsic complexity that is not present in the general coverage problem, since it presents a delicate tradeoff between minimizing the accumulated risk and minimizing the total coverage time. Trying to minimize the risk involved in the coverage path could mean making some redundant steps, which in turn can make the coverage path longer, and thus increase the risk involved, as well as increase the coverage time.

The adversarial coverage problem has many different variants, that depend on the information given to the robot prior to the coverage (offline vs. online coverage), the representation of the environment, the impact of the threats on the covering robot (whether they stop it completely or only delay it for a certain amount of time), whether the threats can change over time, how many robots are used for the coverage, and more.

2. PROBLEM DEFINITION

We are given a map of a target area T, which contains obstacles and also points with threats, which may stop the robot. We assume that T can be decomposed into a regular square grid with n cells, whose size equals the size of the robot. Some cells in T contain threat points. Each threat point i is associated with a threat probability p_i, which measures the likelihood that the threat will stop the robot. The robot's task is to plan a path through T such that every accessible free cell in T is visited by the robot at least once.

Figure 1 shows an example map of the world. Obstacles are represented by black cells, safe cells are colored white and dangerous cells are represented by 5 different shades of purple. Darker shades represent higher values of p_i (more dangerous areas).

We consider two objectives in regard to the robot's survivability:

1. Minimize the total accumulated risk along the coverage path (i.e., maximize the probability of covering the whole target area).
2. Maximize the coverage percentage of the target area before the robot is first hit (i.e., maximize the expected coverage percentage).

Note that for the first objective, the order of visits of the cells is not important, as long as the number of visits of threat points along the coverage path is minimized (ideally,

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

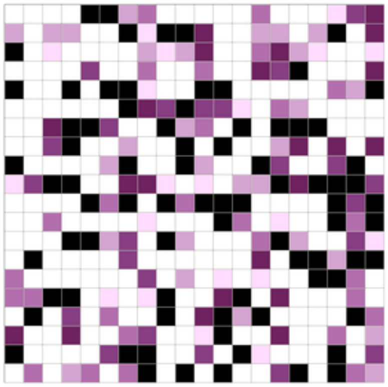

Figure 1: An example map of the world. Darker purple cells represent more dangerous areas.

visiting each threat point only once). On the other hand, for the second objective, the visit order of the cells is crucial, since the robot is trying to cover as much as possible before getting hit by a threat (ideally, covering all the safe cells before visiting a single threat point).

3. MAIN RESULTS

In [3] we have formally defined the offline adversarial coverage problem for a single robot. We have proposed an initial heuristic algorithm that generates a coverage path which tries to minimize a cost function, that takes into account both the survivability of the robot and the coverage path length. However, the heuristic algorithm worked only for obstacle-free areas, and without any guarantees.

In [4] we have addressed a specific version of the adversarial coverage problem, namely, finding the safest coverage path. We have shown that the problem is $\mathcal{NP}$-Complete, and thus we have suggested two heuristic algorithms for solving the safest path problem: STAC and GSAC. STAC (Spanning-Tree Adversarial Coverage) splits the target area into connected areas of safe and dangerous cells, and then it covers the safe areas before moving to the dangerous ones. On the other hand, GSAC (Greedy Safest Adversarial Coverage) follows a greedy approach, which leads the robot from its current location to the nearest safest location which has not been covered yet. We have provided optimality bounds on both algorithms, and proven that these algorithms produce close to optimal solutions in polynomial time. Experimental results have shown that while STAC tends to achieve higher expected coverage, GSAC produces shorter coverage paths with lower accumulated risk.

In [5] we have shown how to model the adversarial coverage problem as a Markov Decision Process (MDP), and proven that finding an optimal policy of the MDP also provides an optimal solution to this problem. Since the state space of the MDP is exponential in the size of the target area's map, we have used real-time dynamic programming (RTDP), a well-known heuristic search algorithm for solving MDPs with large state spaces. Although RTDP achieves faster convergence than value iteration on this problem, practically it cannot handle maps with sizes larger than 7×7. Hence, we have introduced the use of frontiers, states that separate the covered regions in the search space from those uncovered, into RTDP. We have shown that Frontier-Based RTDP (FBRTDP) converges orders of magnitude faster than RTDP, and obtains significant improvement over the state-of-the-art solution for the adversarial coverage problem.

In [2] we have built a more sophisticated model of the adversary, in which it can choose the best locations of the threat points, such that the probability of stopping the covering robot is maximized. In other words, we have examined the problem of finding the best strategy to defend a given area from being covered by an agent, using k given guards. We have examined the impact of the adversarial knowledge of the coverage path on the choice of the guards' locations, and provided solutions for adversaries having no knowledge and full knowledge of the coverage path. We have shown that for a full-knowledge adversary there is a simple algorithm that provides the optimal strategy, whereas finding an optimal strategy for a zero-knowledge adversary is, in general, $\mathcal{NP}$-Hard. However, for some values of k such an optimal strategy can be found in polynomial time, and for others we have suggested heuristics that can significantly improve the random baseline strategy. We have also discussed some cases in which the adversary has partial knowledge of the coverage path (for example, when it only knows where the coverage begins).

4. FUTURE WORK

There are several areas we plan to pursue in future work. First, we are interested in finding algorithms for the online version of the adversarial coverage problem, in which the coverage has to be completed without the use of a map or any a-priori knowledge of the target area. Second, we would like to consider non-stationary environments, where the locations of the threat points can change over time. Finally, we would like to extend the suggested algorithms for multi-robot systems. Using multiple robots for coverage has the potential for more efficient coverage and greater robustness; even if one robot is totally damaged, others may take over its coverage subtask.

REFERENCES

[1] E. Galceran and M. Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12):1258–1276, 2013.

[2] R. Yehoshua and N. Agmon. Adversarial modeling in the robotic coverage problem. In *International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-15)*, 2015.

[3] R. Yehoshua, N. Agmon, and G. A. Kaminka. Robotic adversarial coverage: Introduction and preliminary results. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-13)*, pages 6000–6005, 2013.

[4] R. Yehoshua, N. Agmon, and G. A. Kaminka. Safest path adversarial coverage. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-14)*, pages 3027–3032, 2014.

[5] R. Yehoshua, N. Agmon, and G. A. Kaminka. Frontier-based RTDP: A new approach to solving the robotic adversarial coverage problem. In *International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-15)*, 2015.