

Multi-Robot Adversarial Coverage

Roi Yehoshua and Noa Agmon

Bar Ilan University

{yehoshr1,agmon}@cs.biu.ac.il

Abstract

This work discusses the problem of adversarial coverage, in which one or more robots are required to visit every point of a given area, which contains threats that might stop the robots. Previous studies of the problem dealt with single-robot coverage. Using a multi-robot team for the coverage has clear advantages in terms of both coverage time and robustness: even if one robot is totally damaged, others may take over its coverage subtask. Hence, in this paper we describe a multi-robot coverage algorithm that tries to maximize the percentage of covered area before the team is stopped, while minimizing the coverage time. We analytically show that the algorithm is robust, in that as long as a single robot is able to move, the coverage will be completed. We also establish theoretical bounds on the minimum covered area guaranteed by the algorithm and on the coverage time. Lastly, we evaluate the effectiveness of the algorithm in an extensive set of environments and settings.

1 Introduction

Coverage path planning is one of the fundamental problems in robotics. The goal of coverage path planning is to find a sequence of world locations which allows the robot(s) to visit every part of the target area while optimizing some criteria, usually minimizing travel cost while avoiding obstacles. This problem has many real-world applications, from automatic floor cleaning [Colegrave and Branch, 1994], to field demining [Nicoud and Habib, 1995] and surveillance by unmanned aerial vehicles (UAVs) [Girard *et al.*, 2004].

In a recently introduced version of the problem, *adversarial coverage* (e.g., [Yehoshua *et al.*, 2016]), the robot has to cover the given terrain without being stopped by an adversary. Each point in the area is associated with a probability of the robot being stopped at that point. The objective of the robot is to cover the entire target area (including the threat points) as quickly as possible while minimizing the probability that it will be stopped before completing the coverage. This problem is a generalization of the original problem of coverage in neutral environments, where risks do not exist, thus are not

accounted for, and the only goal is to minimize coverage time (e.g. [Gabriely and Rimon, 2003]).

Previous work of adversarial coverage dealt with the single-robot version of the problem. There are obvious advantages in using multiple robots in the adversarial coverage task. Using multiple robots clearly decreases the time to complete the task due to workload division, as seen in the original multi-robot coverage problem [Hazon and Kaminka, 2008], [Agmon *et al.*, 2006]. Additionally, using multiple robots improves robustness, as failure of some members of the robot team can be compensated by others. Therefore, in this paper we extend adversarial coverage to multi-robot systems. We focus on coverage using a map of the work-area (known as offline coverage [Choset, 2001]).

First, we formally define the multi-robot version of adversarial coverage, and the problem of finding the safest coverage path for a multi-robot team. Second, we describe an efficient distributed multi-robot adversarial coverage algorithm, that tries to maximize the survivability of the team while optimizing the coverage time. The algorithm is based on decomposition of the target area into connected areas of safe and dangerous locations, and prioritization of the coverage such that coverage of safer areas comes before coverage of more dangerous ones. Our method also utilizes graph partitioning techniques in cases where more than one robot is assigned to the coverage of a given area, in order to speed up its coverage. We provide a theoretical bound on the expected coverage percentage guaranteed by the algorithm, and also analyze the best-case and worst-case completion times for the algorithm. Finally, we evaluate our method in an extensive set of experiments. The results show that adding more robots to the coverage task can significantly increase the percentage of the area covered as well as reduce the coverage time.

2 Related Work

The problem of single and multi-robot coverage has been extensively discussed in the robotic literature (see Galceran and Carreras [2013] for a recent survey). Most approaches to multi-robot coverage extend single-robot ideas to multiple robots by using a strategy to divide the workload. Hazon and Kaminka [2008] generalized the STC method to multi-robot teams in the family of Multi Robot Spanning Tree Coverage (MSTC) algorithms. Their solution, along with decreasing the total coverage time, achieved robustness in the sense that

as long as one robot works properly, the coverage of the terrain is guaranteed. They have also shown that in multi-robot teams redundancy might be necessary for more efficiency.

Rekleitis et al. [2008] presented a collection of algorithms for the coverage planning problem using a team of mobile robots on an unknown environment, based on an exact cellular decomposition. To achieve coverage in line-of-sight-only communications, the robots take two roles: some members, called explorers, cover the boundaries of the current target cell, while the other members, called coverers, perform simple back-and-forth motions to cover the remainder of the cell. For task/cell allocation among the robots, a greedy auction mechanism is used.

The offline single-robot adversarial coverage problem was formally defined in [Yehoshua et al., 2013], in which we proposed a simple heuristic algorithm for generating a coverage path aiming at minimizing a cost composed of both the survivability of the robot and the coverage path length. The heuristic algorithm worked only for obstacle-free areas, and without any guarantees. In a follow-up paper [Yehoshua et al., 2014] we have addressed a more specific version of the problem, namely, finding the safest coverage path. There we suggested two heuristic algorithms: STAC, a spanning-tree based coverage algorithm, and GAC, which follows a greedy approach. In [Yehoshua and Agmon, 2015a] 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. Lastly, the online single-robot version of the problem was presented in [Yehoshua and Agmon, 2015b].

3 Problem Formulation

A team of k robots $\mathcal{R} = \{R_1, \dots, R_k\}$ needs to cover a given area. The area contains threats that may stop the robots, as well as obstacles. In contrast to obstacles which the robots cannot go through, threat locations are places that the robots must visit, but might be stopped at. The robots are given a map of the environment in advance. We assume that communication between the robots is available without any restrictions. Each robot autonomously covers the area it is assigned, keeping track of all the covered and uncovered space by communicating with the other robots.

Furthermore, we assume that the given area can be decomposed into a regular grid with n cells. The grid contains two types of cells: free cells and cells that are occupied by obstacles. Some of the free cells contain threats. Each free cell i is associated with a threat probability p_i , which measures the likelihood that a threat in that cell will stop a robot visiting it. The robots can move in the four basic directions (North, South, East, West), and can locate themselves within the work-area to a specific cell.

We assume that robots that have been stopped by threats do not block live robots, i.e. other robots can move through their cells, and that the threats remain there. An example for a real-world scenario in which this assumption holds is when the covering robots are UAVs, and the threats are constant in time (e.g., an anti-drone weapon aimed at a particular location).

The survivability measure can be defined in two levels: the

survivability of each single robot, and the survivability of the team. We say that a robot survives the coverage, if it manages to finish its coverage task unharmed. To formally define the survivability of a robot R_i , we first denote the coverage path that it follows by $P_i = (c_1^i, c_2^i, \dots, c_{n_i}^i)$, where c_j^i is the cell that robot i is located at in time step j . Thus, the probability that a robot R_i survives the coverage is:

$$Surv(R_i) = \prod_{j=1}^{n_i} (1 - p_j^i) \quad (1)$$

where p_j^i denotes the probability that the robot is stopped by a threat in cell c_j^i .

We say that a team of robots survives the coverage, if at least one of the robots in the team survives until all cells in the area are visited (coverage completed). Thus, if it takes t time steps to cover the area, the probability that a team of k robots $\mathcal{R} = \{R_1, \dots, R_k\}$ is able to cover this area is:

$$Surv(\mathcal{R}) = \prod_{j=1}^t \left[1 - \prod_{i=1}^k p_j^i \right] \quad (2)$$

Following these definitions, we can now define the Multi-Robot Safe Adversarial Coverage Problem as follows:

Definition 1. *Multi-Robot Safe Adversarial Coverage Problem (MRSACP):* Given a team of k robots and a grid representation of a world that contains obstacles and threat points, find a coverage path of the grid that maximizes the survivability of the team.

The \mathcal{NP} -hardness of MRSACP follows directly from the \mathcal{NP} -hardness of the single-robot safest coverage path problem [Yehoshua et al., 2014].

4 Multi-Robot Adversarial Coverage

The Multi-Robot Adversarial Coverage algorithm (MRAC, described in Algorithm 1) uses a layered-based approach. It first tries to cover all the safe cells in the target area as efficiently as possible, using all the robots. Then it covers the dangerous areas from the least dangerous ones to the most dangerous ones. This way the algorithm tries to maximize the covered area before the entire team of robots is stopped. We first introduce the following definitions:

Definition 2. A *connected area* is a connected subset of cells in the grid that belong to the same threat level.

Figure 1 shows an example of a grid containing two safe areas and five dangerous areas that belong to two different threat levels. The two low-threat-level areas are outlined by yellow lines and the three high-threat-level areas are outlined by blue lines.

In addition, we denote the distinct threat probabilities that exist in the grid by $p_0, \dots, p_l (p_0 = 0)$ and define *threat level i* as the group of all the cells that contain threats with probability p_i of stopping the robot.

The main steps of the multi-robot coverage algorithm are:

1. Split the target area into layers according to the threat levels, i.e. layer i contains all the cells that belong to threat level i .

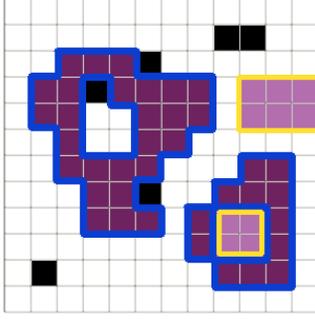


Figure 1: An example of a grid containing two safe areas and five dangerous areas.

2. For each threat level i , denote by $A_1^i, A_2^i, \dots, A_{n_i}^i$ the connected areas that belong to level i .
3. Each robot is assigned to a safe area $A \in \{A_1^0, \dots, A_{n_0}^0\}$, which has the safest path from its current location. Areas with more than one robot assigned to them are split between their assigned robots.
4. For the coverage of each area, we use a modified version of GAC (Greedy Adversarial Coverage), the state-of-the-art solution to the single-robot adversarial coverage problem [Yehoshua *et al.*, 2016].
5. When a robot R_i completes its current coverage task, it is allocated an uncovered area from the safest threat level which has not been completed yet. If all areas have been allocated, then the robot is added to an area that is already being covered by another robot.
6. If a robot is hit by a threat during the coverage of its allocated area, then this area is returned to the pool of unassigned areas.
7. If all the robots have completed their area coverage, and there are no uncovered cells, then the robots declare the environment covered.

4.1 Multi-level Graph Partitioning

There are several places in the algorithm where we need to use a graph partitioning scheme, in order to divide a given area between several robots. The problem of partitioning a graph into k equally-sized components is known to be \mathcal{NP} -Hard [Krauthgamer *et al.*, 2009]. Therefore, practical solutions are based on heuristics. Here we use a multi-level graph partitioning algorithm [Karypis and Kumar, 1998], that consists of three main phases: coarsening, partitioning, and uncoarsening. In the coarsening phase, a sequence of smaller graphs, each with fewer vertices is obtained by collapsing vertices and edges into single vertices of the next level, which are called multi-nodes. Then, in the partitioning phase, the coarse graph obtained is partitioned. Lastly, in the uncoarsening phase, the partitioning is refined while the original graph is restored.

4.2 Data Structures

The algorithm maintains a list of connected areas, denoted by \mathcal{A} . Each area can be in one of the following states:

1. $S_{unassigned}$ - unassigned to any robot
2. $S_{assigned}$ - assigned to a robot and not completely covered yet
3. $S_{covered}$ - covered

For each area A , we maintain the following fields:

- $A.level$ - the threat level this area belongs to
- $A.cells$ - the group of cells that belong to this area
- $A.state$ - the current state of the area
- $A.robot$ - the robot currently assigned to this area

Each robot in the team can be in one of the following states:

1. S_{idle} - waiting for a task assignment. This is the initial state of each robot.
2. $S_{traveling}$ - traveling to the next area that the robot has to cover.
3. $S_{covering}$ - in the process of covering a given area.
4. S_{done} - the robot has finished covering its allocated area and there are no more areas to be covered.
5. S_{dead} - the robot has been hit by a threat.

The transitions between the robot's states are described in section 4.4. For each robot, we maintain the following fields:

- $R.state$ - the current state of the robot
- $R.area$ - the current area the robot is assigned to
- $R.location$ - the robot's location
- $R.path$ - the path the robot is currently following.

In addition to the global map, we assume that the following data structures are shared between the robots:

1. The status of each cell in the map (visited or not).
2. The list of connected areas and their states.
3. The states of all the robots, and the areas that they are assigned to.

4.3 Initial Allocation

Algorithm 1 is responsible for pre-processing of the map and the initial allocation of coverage tasks to the robots. It can be performed offline prior to the coverage, or online by one of the robots, and then its results can be transmitted to all the other robots in the team.

Algorithm 1 Multi-Robot Adversarial Coverage (MRAC)

input: G - the given grid, \mathcal{R} - group of k robots

output: \mathcal{A} - list of connected areas

- 1: $\mathcal{A} \leftarrow \emptyset$
 - 2: Group the cells in G into $l + 1$ threat levels, T_0, \dots, T_l
 - 3: **for each** threat level i , $0 \leq i \leq l$ **do**
 - 4: Build the graph H_i induced from the cells in T_i
 - 5: Find the connected components (areas) of H_i using DFS
 - 6: Let A_1, \dots, A_k be the connected areas of H_i
 - 7: **for each** area A_j , $1 \leq j \leq k$ **do**
 - 8: $A_j.state \leftarrow S_{unassigned}$
 - 9: $A_j.level \leftarrow i$
 - 10: Add A_j to \mathcal{A}
 - 11: Allocate_Areas_To_Robots(\mathcal{A} , \mathcal{R})
-

The procedure `Allocate_Areas_To_Robots`¹ allocates the initial areas to the robots. The idea is to initially cover all the safe areas as fast as possible, using all the robots available, before moving to the dangerous areas. The procedure first computes the safest paths of each robot to every safe area. Then, it assigns each robot to the safe area with the safest path from its current location, if this area is not too dense with robots. We define a *dense* area as an area whose number of cells is less than d times the number of robots assigned to it, i.e., each robot has less than d cells to cover on average. If the area chosen for the robot is already too dense with robots, then the next safest non-dense area will be allocated to it. After the allocation of areas for all the robots, every area that is assigned to more than one robot is split between the robots using the graph partitioning method described in section 4.1.

4.4 The Coverage Algorithm

Algorithm 2 describes the action taken by each robot during one time step during the coverage. The robot typically starts in the state $S_{traveling}$, unless there were not enough areas to allocate to all the robots in the initial phase (e.g., there was only one area with 20 cells and there are 10 robots). In this case, the robot starts in S_{idle} and waits for a task assignment. When the robot arrives at the area that it needs to cover, its state changes to $S_{covering}$. If along the way to its designated area, the robot completes a coverage of another area (e.g., an area that consists of only one cell that resides on the connecting path between two larger areas), then this area’s state is changed to $S_{completed}$, and is removed from the pool of available areas for coverage. Only when the robot arrives at its designated area, the coverage path of this area is computed. This is because by the time the robot gets to this area, some of its cells may have already been visited along the connecting paths of other robots.

When the robot finishes its coverage task, it is assigned a new area to cover from the pool of unassigned areas and its state changes back to $S_{traveling}$. If there are no more unassigned areas, it joins another covering robot to help it finish its coverage task. If there are no more areas that can be shared, then the robot’s state changes to S_{done} and it waits until one of the areas becomes unassigned (e.g., when another robot is stopped). If the robot is stopped during the coverage of its assigned area or on its way to it, its state changes to S_{dead} and its allocated area becomes unassigned.

The procedure `Allocate_Next_Area` allocates a new area to cover for a robot that has completed its coverage task. If there are any unassigned areas, the robot is assigned to the safest unassigned area with the safest path from its current location. If all the areas are already assigned, the robot joins another covering robot to help it finish its coverage task. If an area that can be shared has been found, its connected unvisited parts are defined as new areas and added to the pool of unassigned areas instead of the original area. If there is only one such part (e.g., its assigned robot has not started covering it), then it is split into two balanced parts. Finally, both the assigned robot and the idle robot are (re-)assigned to the

¹Some of the proofs and the pseudocode are omitted due to space constraints. They can be found at <http://goo.gl/ITQk4y>

Algorithm 2 Robot_Action(R)

```

1: switch  $R.state$  do
2:   case  $S_{traveling}$ 
3:      $c \leftarrow$  next cell on  $R.path$ 
4:     Mark  $c$  as visited
5:      $A \leftarrow$  the area that contains  $c$ 
6:     if all cells in  $A$  are visited then
7:        $A.state \leftarrow S_{completed}$ 
8:     if robot was hit by a threat in  $c$  then
9:        $R.state \leftarrow S_{dead}$ 
10:      Reallocate_Area( $R.area$ )
11:     else if  $c$  is the last cell on  $R.path$  then
12:        $R.path \leftarrow$  Area_Coverage( $R.area, R.location$ )
13:        $R.state \leftarrow S_{covering}$ 
14:   case  $S_{covering}$ 
15:      $c \leftarrow$  next cell on  $R.path$ 
16:     Mark  $c$  as visited
17:     if robot was hit by a threat in  $c$  then
18:        $R.state \leftarrow S_{dead}$ 
19:       Reallocate_Area( $R.area$ )
20:     else if  $c$  is the last cell on  $R.path$  then
21:        $R.area.state \leftarrow S_{completed}$ 
22:       Allocate_Next_Area( $R$ )
23:   case  $S_{done}, S_{idle}$ 
24:     if an unassigned area exists in  $\mathcal{A}$  then
25:       Allocate_Next_Area( $R$ )
26:   case  $S_{dead}$ 
27:      $R.area.state \leftarrow S_{unassigned}$ 

```

sub-areas with the safest paths from their current location.

For the coverage of each area, our algorithm is based on the Greedy Adversarial Coverage (GAC), the state-of-the-art solution to the single-robot adversarial coverage problem described in [Yehoshua *et al.*, 2016]. GAC follows a greedy approach, where in each step it leads the robot to the safest nearest cell to its current location which has not been covered yet. We have modified the original coverage algorithm to take into account cells that have already been visited in the target area, since by the time the robot reaches its designated area, other robots may already have visited some cells in this area on the way to their own designated areas.

Finally, the procedure `Reallocate_Area` is used to reallocate an area whose coverage was stopped in the middle. It finds all the unvisited parts of the given area and creates new unassigned areas from them, which are then added to the list of connected areas instead of the given area. The next idle robot will be assigned to one of these areas in its next cycle.

5 Analysis of the MRAC algorithm

We now analyze the MRAC algorithm. We first prove that the algorithm is complete and robust.

Theorem 1. (*Completeness*) *Algorithm MRAC generates paths for the robots that together cover every cell accessible from their starting locations.*

Proof. MRAC partitions the target area into k connected areas, whose union is equal to the target area. Each of these areas is eventually assigned to one of the robots, since no robot remains idle while there are more areas to be covered (lines 24–25 in algorithm 2). The areas are covered by using the

GAC algorithm. Previous work has shown that GAC is complete, i.e. that it produces a path that covers all the accessible cells in its given area (Theorem 8 in [Yehoshua *et al.*, 2016]). Thus, the union of the coverage paths of the k connected areas covers every accessible cell in the target area. \square

Theorem 2. (Robustness) *Algorithm MRAC guarantees that the coverage will be completed as long as at least one robot remains active.*

Proof. The target area is divided into a set of connected areas. Each area is eventually assigned to one of the robots. If a robot is stopped while covering its assigned area, the uncovered parts of this area are returned to the pool of unassigned areas (in procedure `ReallocateArea`). These sub-areas are eventually assigned to one of the remaining robots, since no robot remains idle while there are more areas to be covered (lines 24–25 in algorithm 2). \square

We now provide a bound on the minimum expected number of cells that the robots will cover when following the policy generated by the MRAC algorithm. We will use the definition of expected coverage from [Yehoshua *et al.*, 2016]. Given a coverage path A of a robot R , we denote the sequence of new cells discovered along A by (b_1, \dots, b_n) , and the number of new cells visited by the robot until it is stopped by C_A . Furthermore, for each cell in the sequence b_i , we will denote the sub-path in A that leads from the origin cell a_1 to it by g_i . Then, under the threat probability function p , the expected number of new cells that robot R visits can be expressed as:

$$E(R) = \sum_{i \in (b_1, \dots, b_n)} \prod_{j \in g_i} (1 - p_j) \quad (3)$$

Theorem 7 in [Yehoshua *et al.*, 2016] provides a bound on the minimum expected number of cells that a single robot will be able to cover. Let $A_{i,1}, \dots, A_{i,k_i}$ be the connected areas of threat level i , arranged in the order of their visit by the robot R . Let $|A_{i,j}|$ be the size of area $A_{i,j}$ and $m_{i,j}$ the number of cell visits needed to cover this area. Let $C_{i,j}$ be the set of cells on the connecting path between two consecutive areas $A_{i,j}$ and $A_{i,j+1}$, and C_{i,k_i} be the set of cells on the connecting path between the last area of threat level i and the first area of threat level $i+1$. Denote by $P(C_{i,j})$ the probability to traverse the cells in $C_{i,j}$ without being hit by a threat. Then the expected number of cells robot R will be able to cover before it is stopped is at least:

$$E(R) \geq |A_{0,0}| + \sum_{i=0}^l \sum_{j=1}^{k_i} \left[\prod_{x=0}^{i-1} \left[(1 - p_x)^{\sum_{y=1}^{k_x} m_{x,y}} \prod_{y=1}^{k_x} P(C_{x,y}) \right] \cdot (1 - p_i)^{\sum_{y=1}^j m_{i,y}} \prod_{y=1}^{j-1} P(C_{i,y}) \right] |A_{i,j}| \quad (4)$$

We now use this theorem to prove a lower bound on the expected coverage that can be attained by the entire team.

Theorem 3. *Let \mathcal{R} be a group of k robots $\mathcal{R} = \{R_1, \dots, R_k\}$. Let l be the number of dangerous threat levels. Let \mathcal{A} be the group of connected areas and let C be the set of cells on the connecting path between two areas. Then the expected*

number of cells the team \mathcal{R} will be able to cover before all robots in \mathcal{R} are stopped is at least:

$$E(\mathcal{R}) \geq \sum_{i=1}^k E(R_i) - |C| \quad (5)$$

Proof. Since different robots in the team cover different areas (when one area is assigned to more than one robot, it is split between them), the expected number of cells that will be covered by the entire team is equal to the total number of cells that will be covered by each robot individually minus the number of cells on the connecting paths between areas, since those might be visited multiple times by different robots. \square

Notice that the ideal expected coverage is attained when the robots visit the cells precisely in increasing order of their threat probabilities, i.e., when they first visit all the safe cells, then all the cells with threat probability p_1 , etc. Thus, if we denote the cells that belong to threat level i by $c_{i,1}, \dots, c_{i,n_i}$, then the ideal expected coverage is:

$$\prod_{i=0}^l \sum_{j=1}^{n_i} \left[\prod_{x=0}^{i-1} (1 - p_x)^{n_x} \right] \cdot (1 - p_i)^{j-1} \quad (6)$$

By comparing equations (5) and (6), we can see that the gap between the ideal expected coverage and the expected coverage achieved by MRAC depends on the quality of the coverage algorithm of each connected area. Finding an optimal solution to the coverage problem is \mathcal{NP} -Hard. However, MRAC uses for the coverage of each area the algorithm GAC, which creates a close-to-optimal coverage paths [Yehoshua *et al.*, 2016].

The next theorem provides a bound on the coverage time.

Theorem 4. *The worst-case coverage time for k robots is equal to that of a single robot and the best-case coverage time is approximately $1/k$ the coverage time of a single robot.*

6 Experimental Results

We have fully implemented the MRAC algorithm and evaluated it in various types of simulated environments with varying parameters, such as map size, number of robots, number of threat levels, distribution of threats, etc. We provide here the most interesting results from our experiments.

First, we demonstrate the algorithm's results on a specific grid map of size 20×20 . Figure 2 shows the locations of the robots after the coverage has finished, and the total number of times each cell in the grid has been visited. Obstacles are represented by black cells, safe cells are colored white, and dangerous cells are represented by 5 different shades of magenta. Darker shades represent more dangerous areas. In this example, the maps contained 10 contiguous threat areas which were divided into 5 different threat levels with the following threat probabilities: 4%, 8%, 12%, 16%, and 20%. We ran the algorithm on this map with two team sizes of 4 robots and 10 robots. All the robots start in the upper-left corner of the area. The figure shows the locations of the robots after the coverage has finished, and the total number of times each cell in the grid has been visited by the robots. Using 10 robots instead of 4 robots has increased the coverage percentage from 80.33% to 97.33% and reduced the coverage time from 311 to 244 time steps. As can be seen in both maps, more dangerous areas are less frequently visited by the robots.

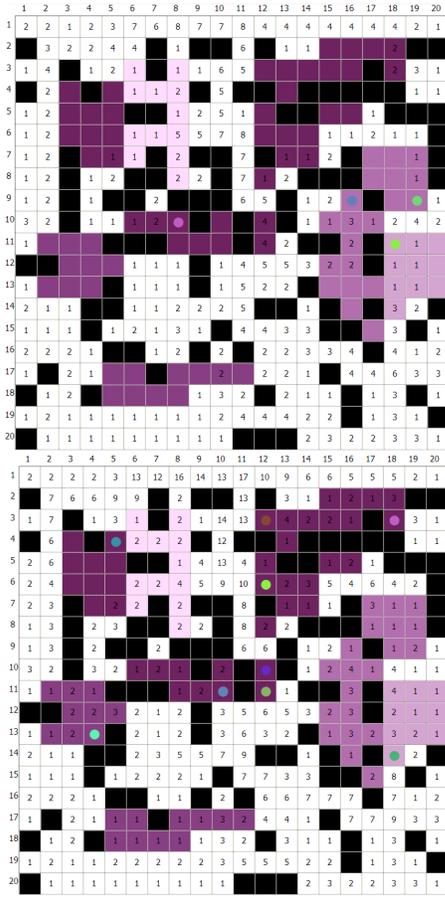


Figure 2: A sample map with the results of MRAC with 4 robots (upper figure) and 10 robots (lower figure). The number in each cell indicates the number of visits.

Second, we examined the effect of changing the team size. We used 500 randomly-generated maps of size 20×20 , with varying number of robots between 1 and 20. In all the experiments, the obstacles ratio and the threats ratio were both set to 25% of the cells. The maps contain 10 contiguous threat areas which are divided into 5 different threat levels with threat probabilities between 4% and 20%. Figure 3 shows the coverage percentage and time for each team size.

As can be seen, the coverage percentage increases and the coverage time decreases as we add more robots to the team. Although the area that the robots are able to cover increases as the team grows, they are able to cover it more quickly. The small increase in the coverage time in the transition from one robot to two robots is due to the fact that the area that two robots are able to cover is significantly larger than one robot (62% for two robots vs. 44% for one robot), thus it takes more time until the two robots are stopped despite the division of the workload between them (e.g., one robot may be stopped much earlier than the second one).

Lastly, we have examined the effect of changing the threats ratio in the environment. Figure 4 shows the coverage percentage and coverage time for various threat ratios between 0% and 35%. In all the experiments we used teams of 8 robots

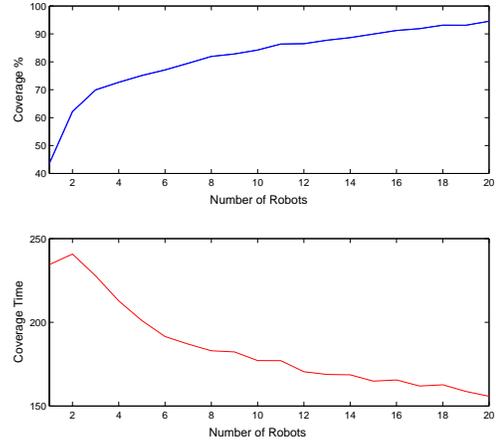


Figure 3: Comparing coverage percentage and time for 1 to 20 robots in environments with contiguous threats.

to cover the area and the threats were concentrated in 10 dangerous areas. As expected, adding more threats the environment causes the team to cover smaller percentage of the area. Note that when the threats ratio is more than 15% the coverage time decreases, since the area that needs to be covered gets smaller thus the team is able to cover it more quickly.

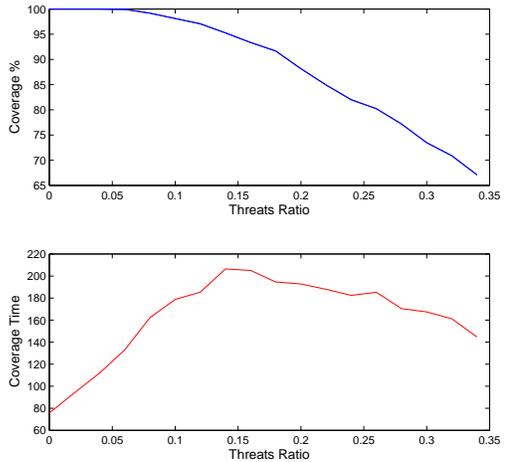


Figure 4: Comparing coverage percentage and coverage time for varying threat ratios between 0 and 0.35.

7 Conclusions and Future Work

We described a multi-robot coverage algorithm for adversarial environments that is complete and robust in face of threats harming the robots. We examined the efficiency of the algorithm in terms of both the survivability of the team and the coverage time, and have shown that adding more robots to the team significantly improves both measures. In future, we would like to extend the multi-robot coverage algorithm to online scenarios, in which the robots are not given a map of the area, and also consider non-stationary environments.

References

- [Agmon *et al.*, 2006] Noa Agmon, Noam Hazon, and Gal A Kaminka. Constructing spanning trees for efficient multi-robot coverage. In *IEEE International Conference on Robotics and Automation (ICRA-06)*, pages 1698–1703, 2006.
- [Choset, 2001] Howie Choset. Coverage for robotics—a survey of recent results. *Annals of mathematics and artificial intelligence*, 31(1-4):113–126, 2001.
- [Colegrave and Branch, 1994] J Colegrave and A Branch. A case study of autonomous household vacuum cleaner. *AIAA/NASA CIRFFSS*, page 107, 1994.
- [Gabriely and Rimon, 2003] Yoav Gabriely and Elon Rimon. Competitive on-line coverage of grid environments by a mobile robot. *Computational Geometry*, 24(3):197–224, 2003.
- [Galceran and Carreras, 2013] Enric Galceran and Marc Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12):1258–1276, 2013.
- [Girard *et al.*, 2004] Anouck R Girard, Adam S Howell, and J Karl Hedrick. Border patrol and surveillance missions using multiple unmanned air vehicles. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 1, pages 620–625, 2004.
- [Hazon and Kaminka, 2008] Noam Hazon and Gal A Kaminka. On redundancy, efficiency, and robustness in coverage for multiple robots. *Robotics and Autonomous Systems*, 56(12):1102–1114, 2008.
- [Karypis and Kumar, 1998] George Karypis and Vipin Kumar. Multilevel k-way partitioning scheme for irregular graphs. *Journal of Parallel and Distributed computing*, 48(1):96–129, 1998.
- [Krauthgamer *et al.*, 2009] Robert Krauthgamer, Joseph Seffi Naor, and Roy Schwartz. Partitioning graphs into balanced components. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 942–949, 2009.
- [Nicoud and Habib, 1995] Jean Daniel Nicoud and Maki K Habib. The Pemex-B autonomous demining robot: perception and navigation strategies. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 'Human Robot Interaction and Cooperative Robots'*, volume 1, pages 419–424, 1995.
- [Rekleitis *et al.*, 2008] Ioannis Rekleitis, Ai Peng New, Edward Samuel Rankin, and Howie Choset. Efficient boustrophedon multi-robot coverage: an algorithmic approach. *Annals of Mathematics and Artificial Intelligence*, 52(2-4):109–142, 2008.
- [Yehoshua and Agmon, 2015a] Roi Yehoshua and Noa Agmon. Adversarial modeling in the robotic coverage problem. In *International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-15)*, pages 891–899, 2015.
- [Yehoshua and Agmon, 2015b] Roi Yehoshua and Noa Agmon. Online robotic adversarial coverage. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-15)*, pages 3830–3835, 2015.
- [Yehoshua *et al.*, 2013] Roi Yehoshua, Noa Agmon, and Gal 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.
- [Yehoshua *et al.*, 2014] Roi Yehoshua, Noa Agmon, and Gal A Kaminka. Safest path adversarial coverage. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-14)*, pages 3027–3032, 2014.
- [Yehoshua *et al.*, 2016] Roi Yehoshua, Noa Agmon, and Gal A Kaminka. Robotic adversarial coverage of known environments. *International Journal of Robotics Research*, Advance online publication, doi:10.1177/0278364915625785, 2016.