

Approximation Algorithms for Distributed Multi-Robot Coverage in Non-Convex Environments

Armin Sadeghi, Ahmad Bilal Asghar and Stephen L. Smith

Department of Electrical and Computer Engineering
University of Waterloo, Waterloo, ON, Canada
{a6sadegh, abasghar, stephen.smith}@uwaterloo.ca

Abstract. In this paper, we revisit the distributed coverage control problem with multiple robots on both metric graphs and in non-convex continuous environments. Traditionally, the solutions provided for this problem converge to a locally optimal solution with no guarantees on the quality of the solution. We consider sub-additive sensing functions, which capture the scenarios where sensing an event requires the robot to visit the event location. For these sensing functions, we provide the first constant factor approximation algorithms for the distributed coverage problem. The approximation results require twice the conventional communication range in the existing coverage algorithms. However, we show through extensive simulation results that the proposed approximation algorithms outperform several existing algorithms in convex, non-convex continuous, and discrete environments even with the conventional communication ranges. Moreover, the proposed algorithms match the state-of-the-art centralized algorithms in the solution quality.

Keywords: Multiple and Distributed Robots, Sensor Networks

1 Introduction

Distributed Coverage is a very well studied problem [10, 12, 21, 22, 25] with extensive multi-robot applications, such as environmental monitoring [11], and surveillance [19]. The objective is to deploy a set of robots to cover an environment such that each robot services or senses the events closer to that robot than any other robot. The events arrive according to a spatial distribution, and the cost of sensing an event is a function of the distance from the robot to that event. The distributed coverage control problem is to minimize the total coverage cost of the environment. The existing distributed algorithms to solve this problem converge to a locally optimal solution with no guarantees on the quality of the solution. In this paper, we provide distributed approximation algorithms to solve the problem in non-convex continuous and discrete environments.

Acknowledgements: This research is partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

The first distributed algorithm for coverage control in convex environments was proposed by Cortes *et al* [10]. The algorithm utilizes Lloyd’s descent to converge to a locally optimal solution, and the partition of each robot is defined using Voronoi partitioning. The robots communicate with the robots in their neighboring partitions to implement the algorithm.

Building on the Lloyd’s descent-based algorithm in [10], there has been extensive studies on the coverage control problem in non-convex environments. In [9, 8], the authors map non-convex environments through a diffeomorphism to a convex region and then solve the problem using the Lloyd’s algorithm [10] before mapping the locally optimal solution back to the original environment. For non-convex polygonal environments, a distributed algorithm was presented in [7], where the Lloyd’s algorithm for convex environments was combined with a local path planning algorithm to avoid obstacles. A key idea introduced in [24] is to consider geodesic distance when computing partitions. In [14] and [18], the authors construct the Voronoi partitions based on the visibility of the robot in the presence of obstacles. Unlike the approaches mentioned above, we discretize the non-convex environment, and solve the coverage problem on the discrete environment and provide guarantees on the solution quality.

The approach of converting a continuous non-convex environment to a discrete environment is used in [12, 2, 6]. We utilize the same approach, but are able to characterize the cost of the solution obtained from the discretized environment in the corresponding continuous environment as a function of the sampling density. The authors in [26] study the coverage problem defined on an undirected graph and present a distributed algorithm that converges to a local optimum. Their algorithm requires the robots to know the information of the neighbors of their neighbors. In this paper, we make the same assumption on the communication range of the robots, but establish approximation guarantees.

A closely related problem to the discretized coverage control is the facility location problem [13, 23] where the objective is to minimize the cost of the robots and the total service time of the demands arriving on the vertices. A special case of this problem is the k -median problem [3, 17, 1] where the objective is to place k robots on vertices of the graph to minimize the total service time. A centralized approximation algorithm was presented for the k -median problem in [3], and the analysis of our distributed approximation algorithm leverages this centralized approximation algorithm. The authors in [4] study a k -median clustering problem with distributed Euclidean data, and provide a randomized constant factor approximation based on distributed coresets computation. In contrast, we consider more general non-convex environments and provide a deterministic approximation algorithm.

Contributions: Our main contributions are threefold. First, given a continuous non-convex coverage problem, we generate a corresponding instance on a metric graph, and characterize the performance of the discrete solution on the continuous problem (Section 3). Second, we provide a constant factor approximation algorithm for the distributed coverage problem on metric graphs (Section 4). To the best of our knowledge, this is the first deterministic approximation algo-

rithm for the distributed coverage problem. We prove the approximation results in Section 5. Third, we show through extensive simulations that the proposed algorithm outperforms several existing approaches in convex and non-convex environments, and matches the centralized algorithms in solution quality (Section 6).

2 Continuous and Discrete Coverage Problems

We begin by reviewing the coverage problems in both continuous [10] and discrete environments [26].

2.1 Continuous Environment

Consider m mobile robots in a compact environment with obstacles and let \mathcal{X} be the obstacle free subset of the environment. There is an event distribution $\phi : \mathcal{X} \rightarrow \mathbb{R}_+$ defined over the environment. Let $d(p, q)$ be the length of the shortest path between two locations p and q in \mathcal{X} . The sensing cost of an event at location p by a robot at q is a strictly non-decreasing function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ of $d(p, q)$. Following the non-convex problem formulation in [7], which extends the original formulation in [10], the continuous problem is defined as the problem of finding the set of locations in the environment for the robots that minimizes the sensing cost of the events, i.e.,

$$\min_{Q \in \mathcal{X}^m} \mathcal{H}(Q) = \min_{Q \in \mathcal{X}^m} \int_{\mathcal{X}} \min_{q_i \in Q} f(d(p, q_i)) \phi(p) dp. \quad (1)$$

Without loss of generality, in the rest of the paper we assume that $\int_{\mathcal{X}} \phi(p) dp = 1$. Observe that the best sensing cost for an event is provided by the closest robot to that event location. Then for a given configuration Q , we partition the environment into Voronoi subsets as follows:

$$V_i(Q) = \{p \in \mathcal{X} \mid d(p, q_i) \leq d(p, q_j) \forall q_j \in Q \setminus \{q_i\}\}.$$

The robots move according to some dynamics $\dot{q}_i = g(q_i, u_i)$ where the computation of the shortest path between two configurations of the robot is tractable. Typically first order dynamics $g(r_i, u_i) = u_i$ is considered for the robots in coverage control literature [10]. We are interested in the distributed version of the coverage problem, where the robots have local information on the other robots and each robot computes its control input locally.

2.2 Discrete Environment

Consider a metric graph $G = (\mathcal{S}, E, c)$ where \mathcal{S} is the vertex set, E is the set of edges between the vertices, c is the metric edge cost and let w be the weight on the vertices. Given a team of m robots, the coverage problem on graph G , is the problem of finding a set of m locations to optimally cover the vertices of G , i.e.,

minimize $\mathcal{D}(Q) = \sum_{v \in \mathcal{S}} \min_{q \in Q} w(v)c(v, q)$. For a given configuration $Q \in \mathcal{S}^m$, we can partition the vertices into m subsets

$$W_i(Q) = \{u \in \mathcal{S} | c(u, q_i) < c(u, q_j) \forall q_j \in Q \setminus \{q_i\}\}.$$

If there exists a vertex that has equal distance to two or more robots in Q , then the vertex is assigned to the robot with smaller unique identifier (UID). Robots travel on the edges of the graph and the control input to a robot is a sequence of edges leading to its destination vertex.

Centralized Approximation Algorithm: The centralized version of this problem is a well-known NP-hard problem called the k -median problem [3]. The best known approximation algorithm for this problem on metric graphs is a centralized local search algorithm which provides solutions within a constant factor of the optimal [1]. Starting from a configuration Q , the centralized local search algorithm swaps p vertices in Q at a time with a subset of p vertices in $\mathcal{S} \setminus Q$. If the new configuration improves the coverage by at least some $\epsilon_0 > 0$, then we call this move a valid local move. The procedure terminates if there are no more valid swaps improving the total sensing cost. We will refer to this local search algorithm as CENTRALIZEDALG in the rest of the paper. The solution obtained from CENTRALIZEDALG is within $3+2/p+o(\epsilon_0)$ of the globally optimal solution.

We focus on the distributed version of this problem introduced in [26] where the robots use only the local information to compute their control input. In the following section, we establish the connection between the continuous and discrete coverage problems.

3 From a Continuous to a Discrete Problem

To establish a connection between the coverage problem in continuous and discrete environments, we first convert the continuous coverage problem to a coverage problem in a discrete environment through sampling of the environment.

Let \mathcal{S} be the set of samples of \mathcal{X} with dispersion ζ [15], where ζ is the maximum distance of any point in the environment \mathcal{X} from the closest point in \mathcal{S} , i.e., $\zeta = \max_{p \in \mathcal{X}} \min_{u \in \mathcal{S}} d(p, u)$, (See Figure 1). We construct a metric graph $G = (\mathcal{S}, E, c)$ on sampled locations \mathcal{S} , where E is the edge set and c is a function assigning costs to the edges of the graph. The cost of an edge between two sampled locations $u, v \in \mathcal{S}$ is $c(u, v) = f(d(u, v))$. Let $\sigma(v)$ for $v \in \mathcal{S}$ be the points in \mathcal{X} closer to v than other samples in \mathcal{S} , i.e.,

$$\sigma(v) = \{p \in \mathcal{X} | d(p, v) \leq d(p, u) \forall u \in \mathcal{S} \setminus \{v\}\}.$$

With a slight abuse of notation, let $\sigma^{-1}(p)$ be the closest sample in \mathcal{S} to $p \in \mathcal{X}$. The function $w : \mathcal{S} \rightarrow \mathbb{R}_+$ assigning weights to the vertices of the graph is $w(v) = \int_{p \in \sigma(v)} \phi(p) dp$. We assume the following property on the sensing function.

Assumption 1 (Subadditivity of sensing function). *We assume that the sensing cost function f is a sub-additive function, i.e.,*

$$f(d(p, u) + d(u, v)) \leq f(d(p, u)) + f(d(u, v)).$$

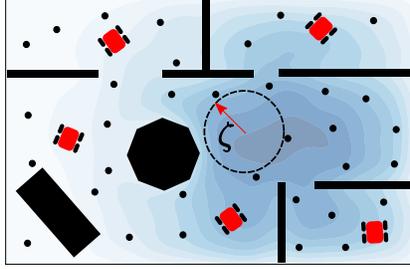


Fig. 1: Sampled locations in an environment with dispersion ζ .

For instance $f(x) = \sqrt{x}$ and $f(x) = x$ are sub-additive functions.

Remark 1. In applications such as dynamic vehicle routing problems [5] and facility location problems [23] where the sensing cost is determined by the distance traveled by the clients, the sensing function falls under Assumption 1.

Due to Assumption 1, the cost function c on the edges of the graph G satisfies the triangle inequality, i.e., for all u, v, z in \mathcal{S}

$$\begin{aligned} c(u, v) &= f(d(u, v)) \leq f(d(u, z) + d(z, v)) \\ &\leq f(d(u, z)) + f(d(z, v)) = c(u, z) + c(z, v). \end{aligned}$$

The following result establishes a connection between the sensing costs of an approximate solution to the discrete coverage problem and the optimal coverage in continuous environment.

Theorem 1. *Consider a continuous coverage problem on environment \mathcal{X} with optimal solution S^* , and its corresponding discrete instance obtained through the set of samples \mathcal{S} with dispersion ζ . Then if Q is the solution obtained from an α -approximation algorithm for the discrete coverage problem instance, the sensing cost of Q on the corresponding continuous problem is $\mathcal{H}(Q) \leq \alpha\mathcal{H}(S^*) + o(f(\zeta))$.*

Proof. We have $\mathcal{H}(Q) = \int_{\mathcal{X}} \min_{q_i \in Q} f(d(q_i, p)) \phi(p) dp$. Therefore,

$$\begin{aligned} \mathcal{H}(Q) &\leq \int_{\mathcal{X}} \min_{q_i \in Q} f(d(q_i, \sigma^{-1}(p)) + d(\sigma^{-1}(p), p)) \phi(p) dp \quad (\text{triangle inequality}) \\ &\leq \int_{\mathcal{X}} \min_{q_i \in Q} [f(d(q_i, \sigma^{-1}(p))) + f(d(\sigma^{-1}(p), p))] \phi(p) dp \quad (\text{Assumption 1}) \\ &= \int_{\mathcal{X}} \min_{q_i \in Q} f(d(q_i, \sigma^{-1}(p))) \phi(p) dp + \int_{\mathcal{X}} f(d(\sigma^{-1}(p), p)) \phi(p) dp \\ &\leq \mathcal{D}(Q) + f(\zeta) \int_{\mathcal{X}} \phi(p) dp = \mathcal{D}(Q) + f(\zeta). \end{aligned} \tag{2}$$

Let $S^* = \{q_i^* | i \in [m]\}$ be the optimal configuration of the continuous problem, and S_G^* be the configuration constructed by moving each robot location in S^*

to the closest sampled location in \mathcal{S} . Also note that,

$$\begin{aligned}
\mathcal{D}(S_G^*) &= \sum_{q_i \in S_G^*} \sum_{u \in W_i} c(u, q_i) w(u) = \sum_{i \in [m]} \sum_{u \in W_i} c(u, q_i) \int_{p \in \sigma(u)} \phi(p) dp \\
&= \sum_{q_i \in S_G^*} \sum_{u \in W_i} \int_{p \in \sigma(u)} f(d(u, q_i)) \phi(p) dp \\
&\leq \sum_{q_i \in S_G^*} \sum_{u \in W_i} \int_{p \in \sigma(u)} [\min_{q_j \in S_G^*} f(d(q_j, p)) + 2f(d(u, p)) + f(d(p, u))] \phi(p) dp \\
&\leq \mathcal{H}(S_G^*) + 3f(\zeta), \tag{3}
\end{aligned}$$

where the first inequality is due to triangle inequality and Assumption 1. Furthermore, we have,

$$\begin{aligned}
\mathcal{H}(S_G^*) &= \sum_{q_i \in S_G^*} \int_{V_i(S_G^*)} f(q_i, p) \phi(p) dp \leq \sum_{q_i \in S_G^*} \int_{V_i(S^*)} f(q_i, p) \phi(p) dp \\
&\leq \sum_{q_i \in S_G^*} \int_{V_i(S^*)} f(d(q_i^*, p)) \phi(p) dp + \int_{\mathcal{X}} f(d(q_i, q_i^*)) \phi(p) dp \\
&\leq \mathcal{H}(S^*) + f(\zeta), \tag{4}
\end{aligned}$$

where the second inequality is due to triangle inequality and Assumption 1. Let Q_G^* be the optimal solution to the discrete coverage problem on graph G , then $D(Q) \leq \alpha \mathcal{D}(Q_G^*) \leq \alpha \mathcal{D}(S_G^*)$. Therefore, by Equations (2), (3) and (4), we have, $\mathcal{H}(Q) \leq \alpha \mathcal{H}(S^*) + (4\alpha + 1)f(\zeta)$. \square

A 5-approximation algorithm for the centralized coverage on metric graphs is provided in [3]. In the following section, we provide the first distributed approximation algorithm for the coverage in metric graphs.

4 Distributed Algorithm On Graphs

In distributed coverage control algorithms for continuous environments, and their adaptations to discrete environments, the algorithm drives each robot to the position inside its partition such that the sensing cost of its partition is minimized, i.e., the centroid of its Voronoi cell in the continuous problem. Although these algorithms converge to locally optimal solutions, there are no global guarantees on the quality of the solution. The following example provides a graph construction where such ‘‘move to centroid’’ algorithms perform poorly.

Example 1. Consider the environment shown in Figure 2 with $3n+1$ vertices, $n+1$ robots and unit costs for the shown edges. We consider the metric completion of the shown graph. The vertices are partitioned into two subsets: 1) V_1 with $2n+1$ vertices and unit weights on the vertices and 2) V_2 with n vertices of weights ϵ for some $0 < \epsilon \ll 1$. The highlighted vertices show the configuration

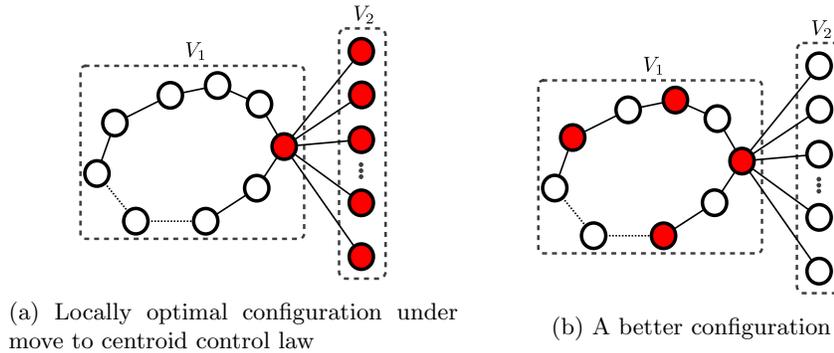


Fig. 2: Example environment with $3n + 1$ vertices and $n + 1$ robots

of the robots. The configuration in Figure 2a is a locally optimal solution under the move to centroid control law with global cost of $n(n + 1)$. However, the configuration shown in Figure 2b provides a global cost of $n + n\epsilon$. Therefore, the locally optimal solution provided by the move to centroid algorithm provides a solution with cost at least $\frac{n+1}{1+\epsilon}$ of the optimal cost on the shown instance.

4.1 High-level Idea

The basic idea of our distributed coverage algorithm is to imitate the local-search algorithm for the k -median problem (See Section 2.2), namely CENTRALIZEDALG, in a distributed manner. The challenge in performing a local move in the distributed manner is that the robots are only aware of the partitions of their neighboring robots, therefore, the effect of a local move on the global objective is not known to the robots. However, we break down a local move in CENTRALIZEDALG into a sequence of moves between neighbors. Let robot j with position q_j and neighbors $\mathcal{N}(j)$ be the closest robot to vertex v . Then a local move of CENTRALIZEDALG swapping the position q_i of robot i with vertex v is equivalent to a sequence of swaps inside Q between the neighboring robots and a move from q_j to v . Figure 3 shows an example of a local move in the centralized algorithm performed by a sequence of local moves.

For this distributed coverage algorithm, we define the minimum communication range and neighbouring robots as follows:

Definition 1 (Neighbour robots). *Given a configuration Q , the set of neighbours of robot i is defined as*

$$\mathcal{N}(i) = \{j \in [m] \mid d(q_i, q_j) \leq 4 \max\{\max_{p \in V_i} d(p, q_i), \max_{p \in V_j} d(p, q_j)\}\}$$

where V_i is the Voronoi partition of robot i in the continuous environment.

An equivalent definition of the neighbour robots in a discrete problem instance where the underlying continuous instance is unknown is provided in [20].

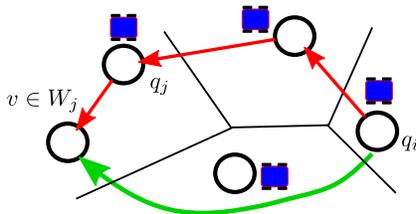


Fig. 3: Local move in the centralized algorithm (green) and its equivalent sequence of moves in the distributed algorithm (red).

Remark 2. The conventional definition of neighbours in the literature [10] is that two robots are neighbours if the intersection of their Voronoi cell boundaries is not empty. Therefore, the distance of two neighbouring robots i and j can be $2 \max\{\max_{p \in V_i} d(q_i, p), \max_{p \in V_j} d(q_j, p)\}$. In [26], authors show that in environments represented as graphs, with the conventional communication range, a move inside a robots partition might change the partition of the neighbours of their neighbours. Therefore, they assume that the robots communicate with the neighbours of their neighbours, which is analogous to twice the communication range needed to implement the Lloyd's descent-based algorithms in continuous environments. In Section 6, we evaluate the performance of the algorithm with both the conventional and our definition of neighbours.

4.2 Detailed Description

We are now ready to provide a detailed description of the proposed algorithm.

Algorithm Framework (Algorithm 1): For the ease of presentation, we provide a description of the algorithm in which the robots perform local moves sequentially. A remark about how the algorithm can be implemented with parallel execution of the local moves is given in [20]. Each robot is assigned a unique identifier UID. Starting from an active robot, say robot i , the robot will make the possible local moves using Algorithm 2. If it can not make a local move, the robot will become inactive and will send a completion message to the neighbouring robots via SENDCOMPLETIONMESSAGE (line 6 of Algorithm 1). After execution of the local move by a robot, the robot becomes inactive. The next active robot to execute the local move can be selected in a distributed manner using a token passing algorithm [16], or any other method that ensures each robot gets a turn at making a local move. The process terminates when all the robots become inactive.

Local Move of Robot i (Algorithm 2): At an iteration of Algorithm 2, let the current configuration of robots be given by $Q = \{q_1, \dots, q_m\}$ where the vertices in the partition of robot i are given by W_i . The robot i considers moving to a vertex $v \in W_i$ from its current vertex q_i . This move can only change the sensing cost of the vertices in the neighbouring robots' partitions (See Lemma 1 in Section 5). Hence, the robot i can calculate the new neighboring partitions

Algorithm 1 DISTRIBUTEDCOVERALGORITHM

```

1: Each robot sets itself to active
2: while there exists an active robot do
3:   for any active robot  $i \in \{1, \dots, m\}$  do
4:     if  $\sum_{u \in W_i} c(u, q_i) > 0$  then
5:       LOCALMOVE( $i$ )
6:       SENDCOMPLETIONMESSAGE()
7:     Robot  $i$  deactivates

```

Algorithm 2 LOCALMOVE(i)

```

1: while  $\exists$  a local move do
2:   Calculate  $\delta_v$  for all vertices in  $W_i$ 
3:   if  $\min_v \delta_v \leq -\epsilon_0$  then
4:     Move to  $v$ 
5:   else
6:     Calculate  $\Gamma_i = \{\rho_v | v \in W_i\}$  ▷ Using Equation (5)
7:     SENDMESSAGE( $\Gamma_i, 1$ )
8:     RECEIVEMESSAGE()
9:     SENDACKNOWLEDGEMENT()

```

after a potential move to v . In line 2 of Algorithm 2, robot i calculates the change in local objective δ_v due to this move for all $v \in W_i$. Since only robot i is executing Algorithm 2 at the current time, δ_v also represents the change in the global objective function. If $\min_{v \in W_i} \delta_v \leq -\epsilon_0$, robot i moves to $q'_i = \arg \min \delta_v$ and the iteration terminates (local move type 1). If there is no valid local move of type 1, then robot i calculates the change in local objective if it moves to v and a new robot appears at q_i , i.e.,

$$\rho_v = \sum_{j \in \mathcal{N}(i)} \sum_{u \in R_j(v)} w(u)[c(u, v) - c(u, q_j)], \quad (5)$$

where $R_j(v) = \{u \in W_j | c(u, q_j) > c(u, v)\}$ represents the vertices in the partition of robot j that are closer to v than the robot j at q_j . Robot i then passes the message with the set $\Gamma_i = \{\rho_v | v \in W_i\}$ and a counter set to 1 to all its neighbors (line 7 of Algorithm 2) and waits for a response (line 8 of Algorithm 2). If the response is a rejection from all the neighbors, Algorithm 2 terminates. Otherwise it selects the acceptance message with the largest change in the objective and moves to the corresponding vertex. It also sends an acknowledgement message to the neighbor k whose message was selected so that robot k can move to q_i .

Response of Other Robots (Algorithm 3): When a robot k receives messages from its neighbors, it follows Algorithm 3. Since messages from only one sender robot are propagating through the system at any time, it can select the message with the smallest counter value if it receives messages from multiple neighbors. The neighbor who sent the message with the smallest counter value is called the parent of robot k . It sends back a rejection message to all other neighbors. If the

Algorithm 3 RECEIVER

Input: message = $(\Gamma_i, \text{MessageCounter})$

- 1: **if** MessageCounter = 1 **then**
- 2: Calculate ℓ_v for all v in the message ▷ Equation (6)
- 3: **if** $\exists v$ in the message with $\rho_v + \ell_v \leq -\epsilon_0$ **then**
- 4: send acceptance message to parent
- 5: **else**
- 6: SENDMESSAGE($\Gamma_i, \text{MessageCounter} + 1$)
- 7: **else**
- 8: Calculate ℓ ▷ Equation (7)
- 9: **if** $\exists v$ in the message with $\rho_v + \ell \leq -\epsilon_0$ **then**
- 10: send acceptance message to parent
- 11: **else**
- 12: SENDMESSAGE($\Gamma_i, \text{MessageCounter} + 1$)
- 13: RECEIVEMESSAGE()

counter value of the message was one, it means that the message originated from its neighboring robot, say i . Then robot k calculates the change in the sensing costs of the vertices in $W_k \setminus R_k(v)$ if it moves to vertex q_i and robot i moves to vertex v resulting in configuration $Q' = \{q_j | j \in \mathcal{N}(k)\} \cup \{v\}$, i.e.,

$$\ell_v = \sum_{u \in W_k \setminus R_k(v)} \min_{q \in Q'} [c(u, q) - c(u, q_k)]w(u), \quad (6)$$

If $\min_v(\rho_v + \ell_v) \leq -\epsilon_0$, robot k decides to move to q_i and sends an acceptance message to robot i with the vertex $\arg \min_v \rho_v + \ell_v$ and the change associated with this move. Otherwise it increments the counter and sends the message with Γ_i to its neighbors. If the counter value in a message is greater than one, robot k calculates ℓ as follows:

$$\ell = \sum_{u \in W_k} \min_{j \in \mathcal{N}(k)} [c(u, q_j) - c(u, q_k)]w(u) \quad (7)$$

If $\min_v(\rho_v + \ell) \leq -\epsilon_0$, robot k sends an acceptance message back to its parent. Otherwise it increments the counter and sends the message to its neighbors.

In function RECEIVEMESSAGE in line 8 of Algorithm 2 and line 13 of Algorithm 3, if any robot receives at least one acceptance message from its neighbors, it passes the message with lowest increase in sensing cost value to its parent. If it receives rejection messages from all its neighbors, it sends back a rejection message to its parent. Robot i selects the move with maximum improvement in the sensing cost and sends back an acknowledgment using SENDACKNOWLEDGMENT to the accepted messages. Robots that receive the acknowledgment move to their parent's location. If there is no more local move available, robot i sends a completion message using SENDCOMPLETIONMESSAGE to the neighbouring robots which will be propagated to all the robots. Then the next active robot executes LOCALMOVE.

5 Analysis of the Algorithm

In this section, we provide analysis on the quality of the solutions provided by the proposed algorithm. The results showing that the proposed algorithm terminates in finite time is provided in [20].

5.1 Correctness and Approximation Factor

Prior to providing the main results on the correctness and approximation factor of the algorithm, we provide two results on the change in the sensing cost of vertices in the partitions of the neighbouring robots with a move of a robot. The proofs of these results are omitted due to space constraints and are given in [20].

The following result shows that a move by a robot inside its partition can only change the sensing cost of the vertices in its neighbouring partitions.

Lemma 1. *Consider a vertex $z \in W_j$ where robot j at position q_j is the closest robot to z . Then robot j is closer to z than any vertex in the partition of a non-neighbour robot, i.e., $c(z, q_j) \leq \min_{i \notin \mathcal{N}(j)} \min_{u \in W_i} c(z, u)$.*

Also, the following result shows that if a robot i moves anywhere in the graph, then the vertices previously in W_i will be assigned to robot in $\mathcal{N}(i)$.

Lemma 2. *For any vertex $z \in W_i$, there exists a robot $j \in \mathcal{N}(i)$ at q_j where $c(z, q_j) \leq c(z, q_k)$ for all $k \notin \mathcal{N}(i)$.*

Then we provide the following result on the change in the global objective with a successful move in the distributed algorithm.

Lemma 3. *If a local move is accepted by the robot, then the global objective improves by at least ϵ_0 .*

Proof. A local move falls under the following cases:

- (i) Since, by Lemma 1, a move of type 1 can only change the sensing cost of the neighbouring robots. Then the result is trivial for the local moves of type 1.
- (ii) If the local move consists of a move by the robot i that is executing LOCAL-MOVE to a vertex v in its partition and a neighbouring robot j moving to vertex q_i . Let Q' be the configuration after the local move, then the change in the global objective $\Delta\mathcal{D} = \mathcal{D}(Q') - \mathcal{D}(Q)$ is given by the following:

$$\Delta\mathcal{D} = \sum_{k \in [m]} \sum_{u \in W_k} \min_{q \in Q'} w(u)c(u, q) - \sum_{k \in [m]} \sum_{u \in W_k} w(u)c(u, q_k). \quad (8)$$

They by Lemma 1 and 2, the sensing cost changes only for vertices in $u \in \cup_{k \in \mathcal{N}(i)} W_k$, therefore,

$$\begin{aligned} \Delta\mathcal{D} &= \sum_{k \in \mathcal{N}(i)} \sum_{u \in W_k} w(u) [\min_{q \in Q'} c(u, q) - c(u, q_k)] \\ &= \sum_{k \in \mathcal{N}(i)} \sum_{u \in R_k(v)} w(u) [c(u, v) - c(u, q_k)] \\ &+ \sum_{k \in \mathcal{N}(i)} \sum_{u \in W_k \setminus R_k(v)} w(u) [\min_{q \in Q'} c(u, q) - c(u, q_k)]. \end{aligned}$$

Observe that the sensing cost for vertex $u \in W_k \setminus R_k(v)$ for robot k at $q_k \in Q' = \mathcal{N}(i) \cup \{v\} \setminus \{j\}$ does not change, therefore, we have

$$\begin{aligned} \Delta \mathcal{D} = & \sum_{k \in \mathcal{N}(i)} \sum_{u \in R_k(v)} w(u)[c(u, v) - c(u, q_k)] \\ & + \sum_{u \in W_j \setminus R_j(v)} w(u)[\min_{q \in Q'} c(u, q) - c(u, q_j)]. \end{aligned}$$

Hence, the result follows immediately as $\Delta \mathcal{D} = \rho_v + l_v$.

- (iii) Suppose the local move consists of a move by the robot i that is executing LOCALMOVE to a vertex v in its partition and a sequence of moves between the neighbouring robots. Without loss of generality, let $\langle v, q_i, q_{i+1}, \dots, q_{j-1}, q_j \rangle$ be the sequence of moves between the neighbouring robots where each robot moves to the previous vertex of the preceding robot in the sequence. Let Q' be the configuration after the local move, then the change in the global objective is given by Equation (8). Observe that each robot accepts only the message from the parent robot. Therefore, among the neighbours of the robots in the sequence only the parent of each robot moves.

First we show that the change in the sensing cost under this sequence of moves only occurs for vertices in $\cup_{k \in \mathcal{N}(i)} W_k$ and vertices in W_j . Let u be a vertex assigned to robot $k' \in [m] \setminus \{j \cup \mathcal{N}(i)\}$ in configuration Q , i.e., $u \in W_{k'}$. Since $k' \notin \mathcal{N}(i)$, then by Lemma 1 a move to vertex v will not improve the sensing cost of u . Also if k' is among the robots moving in the sequence, then there is a robot moving to its previous location, therefore, each vertex in $W_{k'}$ will be sensed by another robot with the same sensing cost. Therefore, the total change $\Delta \mathcal{D}$ in the sensing cost of the vertices becomes

$$\sum_{k \in \mathcal{N}(i)} \sum_{u \in R_k(v)} w(u)[c(u, v) - c(u, q_k)] + \sum_{u \in W_j} w(u)[\min_{q \in Q'} c(u, q) - c(u, q_j)].$$

Hence, the result follows immediately as $\Delta \mathcal{D} = \rho_v + l$. \square

Now we show the following result on the valid local moves in CENTRALIZEDALG given the final configuration of the proposed distributed algorithm.

Lemma 4. *If the proposed distributed algorithm terminates, then there is no single swap move in the centralized local search algorithm CENTRALIZEDALG (See Section 2.2) that improves the objective function.*

Proof. Suppose that there exists a centralized local move of robot j at vertex q_j to a vertex $v \in W_i$ that improves the objective function by ϵ_0 . Therefore, adding a robot to v improves the sensing cost of the vertices in $\cup_{k \in \mathcal{N}(i)} W_k$ by $\rho_v \leq -\epsilon_0$. Therefore, by construction of the algorithm, the robot i would have suggested the move to its neighbouring robots. Suppose after l communications, robot j at q_j receives the message for the first time. In Line 2 (resp. Line 8) of Algorithm 3 if $j \in \mathcal{N}(i)$ (resp. $j \notin \mathcal{N}(i)$), robot j calculates the increase in the sensing cost ℓ_v (resp. ℓ) for the vertices in $\cup_{k \in \mathcal{N}(j)} W_k$ by the move from q_j to the parent of robot j . Since robot j has rejected this offer, by Lemma 3 the change in the global sensing cost is less than ϵ_0 . This is a contradiction. \square

Theorem 2. *The proposed distributed coverage control algorithm provides a solution within $5 + o(\epsilon_0)$ factor of the optimal configuration.*

Proof. The result follows immediately from Lemma 4. The final configuration in the distributed algorithm is a locally optimal solution for the CENTRALIZEDALG with single swap at each iteration, i.e. $p = 1$, therefore, the configuration provides a coverage within a factor $5 + o(\epsilon_0)$ of the global optimal. \square

Corollary 1. *Given an environment \mathcal{X} with m mobile robots and a sampling of \mathcal{X} with dispersion ζ , the solution Q obtained from the proposed distributed coverage control algorithm provides coverage cost $\mathcal{H}(Q) \leq 5\mathcal{H}(S^*) + o(f(\zeta) + \epsilon_0)$, where S^* is the optimal solution of the continuous coverage problem.*

Proof. Proof follows immediately from Theorems 1 and 2. \square

We characterize the run-time of the proposed distributed algorithm to be

$$\log(\mathcal{D}(Q_0)/\mathcal{D}(Q^*))/\log\left(1 - \frac{\epsilon w_0}{p(|\mathcal{S}|, m)} \frac{\min_{e \in E} c(e)}{\max_{e \in E} c(e)}\right),$$

where Q_0 and Q^* are initial and optimal configurations, respectively, $w_0 = \min_{u \in \mathcal{S}} w(u)$, and $p(|\mathcal{S}|, m)$ is a polynomial in $|\mathcal{S}|$ and m . We refer the reader to [20] for the detailed description of the time complexity results.

6 Simulation Results

In this section, we evaluate the performance of the proposed distributed algorithm and compare it to convex and non-convex distributed coverage algorithms and the centralized algorithm in [1]. To construct the discrete problem described in Section 3, we use a grid sampling of the environment. We denote the maximum distance inside a Voronoi cell of a robot i by $R_{\text{comm}} = \max_{p \in V_i} d(q_i, p)$ and we evaluate the performance of the proposed algorithm with communication ranges of $4R_{\text{comm}}$ (See Definition 1) and $2R_{\text{comm}}$ which is analogous to the conventional communication model in the continuous coverage literature. In the rest of this section, we use $f(x) = x$ as the sensing cost function.

6.1 Convex Environments

In this experiment, we compare the proposed algorithm to distributed Lloyd's algorithm [10] in convex environments. We use the Euclidean distance as the metric between two points. The comparison is conducted in a 1500×850 environment with 100 different event distributions. The event distributions are truncated multivariate normal distributions with mean $[1400, 800]$ and covariance matrices $\Sigma = [\sigma, 0; 0, \sigma]$ where σ is uniformly randomly selected from interval $[5, 10] \times 10^4$. In this experiment, the robots are initialized in the bottom left corner of the environment. Figure 4 illustrates the percentage difference of the solutions provided by the two algorithms with respect to the solution of the centralized algorithm. Observe that the proposed achieves solution quality very close to the centralized algorithm, even with a large number of robots, while Lloyd's algorithm provides solutions with approximately 15% deviation.

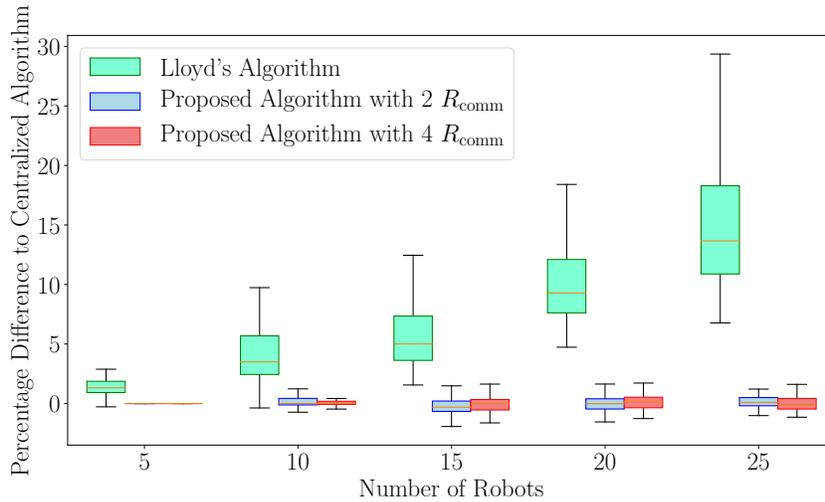


Fig. 4: Comparison of the solutions of different algorithms in a convex environment

6.2 Non-Convex Environments

In this section, we compare the solution quality of the proposed algorithm with two different communication ranges to the algorithms in [7], [26] and the centralized algorithm [1]. The experiment is conducted in a 1500×850 environment that contains obstacles (See Figure 1), and using 100 different event distributions. The distributions are generated in the same manner as in the convex environment experiments with uniformly random mean and covariance matrices. The communication model in the non-convex studies are different, for instance, two robots are neighbours in [7], if the intersection of the Voronoi cells of the robots in the environment without obstacles is non-empty, and two robots are neighbours in [26] if the two partitions of the robots share an edge in the discrete representation of the environment. Therefore in the implementations of algorithms in [7] and [26], we assume that robots are connected to every other robot. Figure 5 shows the percentage difference between the solutions of each algorithm compared to the centralized algorithm. Observe that the proposed algorithm even with the conventional communication range out-performs both other algorithms by $\approx 20\%$ on average in a system with 30 robots and matches the solution quality of the centralized algorithm.

7 Conclusion

This paper considered the multi-robot coverage problem in convex and non-convex environments. A connection is established between the solution quality of the continuous coverage problem and the solution to the coverage problem on a discrete representation of the environment. We also propose the first distributed

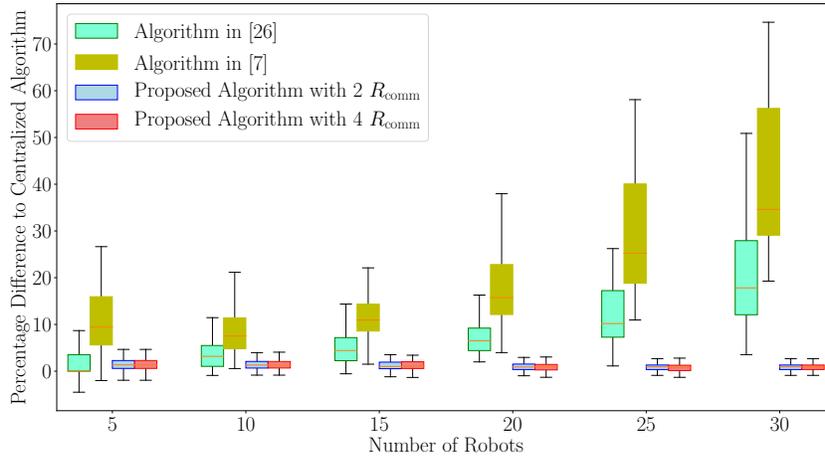


Fig. 5: Comparison of the solutions of different algorithms in a non-convex environment

approximation algorithm for the coverage problem in discrete and continuous environments and provide bound on the quality of the solution. We also characterize the run-time and communication complexity of the proposed algorithm.

References

1. Ahmadian, S., Friggstad, Z., Swamy, C.: Local-search based approximation algorithms for mobile facility location problems. In: Proceedings of the 24th annual ACM-SIAM symposium on Discrete algorithms, pp. 1607–1621. SIAM (2013)
2. Alitappeh, R.J., Pereira, G.A., Araújo, A.R., Pimenta, L.C.: Multi-robot deployment using topological maps. *Journal of Intelligent & Robotic Systems* **86**(3-4), 641–661 (2017)
3. Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., Pandit, V.: Local search heuristics for k -median and facility location problems. *SIAM Journal on computing* **33**(3), 544–562 (2004)
4. Balcan, M.F.F., Ehrlich, S., Liang, Y.: Distributed k -means and k -median clustering on general topologies. In: Advances in Neural Information Processing Systems, pp. 1995–2003 (2013)
5. Bertsimas, D.J., Van Ryzin, G.: A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research* **39**(4), 601–615 (1991)
6. Bhattacharya, S., Michael, N., Kumar, V.: Distributed coverage and exploration in unknown non-convex environments. In: Distributed Autonomous Robotic Systems, pp. 61–75. Springer (2013)
7. Breitenmoser, A., Schwager, M., Metzger, J.C., Siegwart, R., Rus, D.: Voronoi coverage of non-convex environments with a group of networked robots. In: IEEE International Conference on Robotics and Automation, pp. 4982–4989 (2010)
8. Caicedo-Núñez, C.H., Zefran, M.: A coverage algorithm for a class of non-convex regions. In: IEEE Conference on Decision and Control, pp. 4244–4249 (2008)
9. Caicedo-Núñez, C.H., Zefran, M.: Performing coverage on nonconvex domains. In: IEEE International Conference on Control Applications, pp. 1019–1024 (2008)

10. Cortes, J., Martinez, S., Karatas, T., Bullo, F.: Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation* **20**(2), 243–255 (2004)
11. Curtin, T.B., Bellingham, J.G., Catipovic, J., Webb, D.: Autonomous oceanographic sampling networks. *Oceanography* **6**(3), 86–94 (1993)
12. Durham, J.W., Carli, R., Frasca, P., Bullo, F.: Discrete partitioning and coverage control for gossiping robots. *IEEE Transactions on Robotics* **28**(2), 364–378 (2012)
13. Jain, K., Vazirani, V.V.: Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *Journal of the ACM (JACM)* **48**(2), 274–296 (2001)
14. Kantaros, Y., Thanou, M., Tzes, A.: Visibility-oriented coverage control of mobile robotic networks on non-convex regions. In: *IEEE International Conference on Robotics and Automation*, pp. 1126–1131 (2014)
15. LaValle, S.M.: *Planning algorithms*. Cambridge university press (2006)
16. Lemaire, T., Alami, R., Lacroix, S.: A distributed tasks allocation scheme in multi-uav context. In: *IEEE International Conference on Robotics and Automation*, vol. 4, pp. 3622–3627 (2004)
17. Li, S., Svensson, O.: Approximating k-median via pseudo-approximation. *SIAM Journal on Computing* **45**(2), 530–547 (2016)
18. Mahboubi, H., Sharifi, F., Aghdam, A.G., Zhang, Y.: Distributed coordination of multi-agent systems for coverage problem in presence of obstacles. In: *IEEE American Control Conference*, pp. 5252–5257 (2012)
19. Meguerdichian, S., Koushanfar, F., Qu, G., Potkonjak, M.: Exposure in wireless ad-hoc sensor networks. In: *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*, pp. 139–150. ACM (2001)
20. Sadeghi, A., Asghar, A.B., Smith, S.L.: Approximation algorithms for distributed multi-robot coverage in non-convex environments. *arXiv preprint arXiv:2005.02471* (2020)
21. Sadeghi, A., Smith, S.L.: Coverage control for multiple event types with heterogeneous robots. In: *IEEE International Conference on Robotics and Automation*, pp. 3377–3383 (2019)
22. Santos, M., Diaz-Mercado, Y., Egerstedt, M.: Coverage control for multirobot teams with heterogeneous sensing capabilities. *IEEE Robotics and Automation Letters* **3**(2), 919–925 (2018)
23. Shmoys, D.B.: Approximation algorithms for facility location problems. In: *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pp. 27–32. Springer (2000)
24. Thanou, M., Stergiopoulos, Y., Tzes, A.: Distributed coverage using geodesic metric for non-convex environments. In: *IEEE International Conference on Robotics and Automation*, pp. 933–938 (2013)
25. Wang, B.: Coverage problems in sensor networks: A survey. *ACM Computing Surveys (CSUR)* **43**(4), 1–53 (2011)
26. Yun, S.k., Rus, D.: Distributed coverage with mobile robots on a graph: locational optimization and equal-mass partitioning. *Robotica* **32**(2), 257–277 (2014)