Pairwise Optimal Discrete Coverage Control for Gossiping Robots

Joseph W. Durham   Ruggero Carli   Francesco Bullo

Abstract— We propose distributed algorithms to automatically deploy a group of robotic agents and provide coverage of a discretized environment represented by a graph. The classic Lloyd approach to coverage optimization involves separate centering and partitioning steps and converges to the set of centroidal Voronoi partitions. In this work we present a novel graph coverage algorithm which achieves better performance without this separation while requiring only pairwise "gossip" communication between agents. Our new algorithm provably converges to an element of the set of pairwise-optimal partitions, a subset of the set of centroidal Voronoi partitions. We illustrate that this new equilibrium set represents a significant performance improvement through numerical comparisons to existing Lloyd-type methods. Finally, we discuss ways to efficiently do the necessary computations.

I. INTRODUCTION

Coordinated networks of robots are already in use for environmental monitoring [1] and warehouse logistics [2]. In the near future, improvements to the capabilities of autonomous robots will enable robotic teams to revolutionize transportation and delivery of products to customers, search and rescue, and many other applications. All of these tasks share a common feature: the robots are asked to provide service over a space. The distributed *territory partitioning problem* for robotic networks consists of designing individual control and communication laws such that the team will divide a space into territories. Typically, partitioning is done so to optimize a cost function which measures the quality of service provided by the team. *Coverage control* additionally optimizes the positioning of robots inside a territory.

This paper describes a distributed coverage control algorithm for a network of robots to optimize the response time of the team to service requests in an environment represented by a graph. Optimality is defined with reference to a cost function which depends on the locations of the agents and geodesic distances in the graph. As with all multiagent coordination applications, the challenge comes from reducing the communication requirements: the proposed algorithm requires only "gossip" communication, that is, asynchronous and unreliable pairwise communication.

A broad discussion of partitioning and coverage control is presented in [3] which builds on the classic work of Lloyd [4] on algorithms for optimal quantizer design through "centering and partitioning." The relationship between discrete

and continuous coverage control laws based on Euclidean distances is studied in [5]. Coverage control and partitioning of discrete sets are also related to the literature on the facility location or k-center problem [6]. Coverage control algorithms for non-convex environments are discussed in [7], [8], [9] while equitable partitioning is studied in [10]. Other works considering decentralized methods for coverage control include [11] and [12].

In [13] the authors have showed how a group of robotic agents can optimize the partition of convex bounded subset of $\mathbb{R}^d$ using a Lloyd-type algorithm with pairwise "gossip" communication: only one pair of regions is updated at each step of the algorithm. This gossip approach to Lloyd optimization was extended in [14] to discretized non-convex environments more suitable for physical robots.

There are three main contributions of this paper. First, we present a novel gossip coverage algorithm and prove it converges to an element in the set of pairwise-optimal partitions in finite time. This solution set is shown to be a strict subset of the set of centroidal Voronoi partitions, meaning the new algorithm has fewer local minima than Lloyd-type methods. Second, through realistic Player/Stage simulations we show that the set of pairwise-optimal partitions avoids many of the local minima which can trap Lloyd-type algorithms far from the global optimum. Third, we discuss how to efficiently compute our new pairwise coverage optimization.

This paper is organized as follows. Section II defines the domain and goal of our algorithm, while III contains a review of Lloyd-type gossip algorithms from [14]. Section IV presents our new algorithm and its properties, while in V we detail its computational requirements. In Section VI we provide numerical results and comparisons to prior methods, and we end with concluding remarks in VII.

In our notation, $\mathbb{R}_{\geq 0}$ denotes the set of non-negative real numbers and $\mathbb{Z}_{\geq 0}$ the set of non-negative integers. Given a set X, $|X|$ denotes the number of elements in X. Given sets X, Y, their difference is $X \setminus Y = \{x \in X \mid x \notin Y\}$. A set-valued map, denoted by $T : X \rightrightarrows Y$, associates to an element of X a subset of Y.

II. PROBLEM FORMULATION

The problem domain and goal we are considering is the same as that in [14]: given a group of N robotic agents with limited sensing and communication capabilities, and a discretized environment, we want to apportion the environment into smaller regions and assign one region to each agent. The goal is to optimize the quality of coverage, as measured by a cost functional which depends on the current partition and the positions of the agents.

Let the finite set Q be the discretized environment. We assume that the elements of Q, which can be thought of as locations, are connected by weighted edges. Let $G = (Q, E, w)$ be an (undirected) weighted graph with edge set $E \subset Q \times Q$ and weight map $w : E \to \mathbb{R}_{>0}$; we let $w_e > 0$ be the weight of edge e. We assume that G is connected and think of the edge weights as distances between locations.

In any weighted graph G there is a standard notion of distance between vertices defined as follows. A *path* in G is an ordered sequence of vertices such that any consecutive pair is an edge of G. The *weight of a path* is the sum of the weights of the edges in the path. Given vertices h and k in G, the *distance* between h and k, denoted $d_G(h, k)$, is the weight of the lowest weight path between them, or $+\infty$ if there is no path. If G is connected, then the distance between any two vertices is finite. By convention, $d_G(h, k) = 0$ if $h = k$. Note that $d_G(h, k) = d_G(k, h)$, for any $h, k \in Q$.

We will be partitioning Q into N connected subsets or territories to be covered by individual agents. To do so we need to define distances on induced subgraphs of $G = (Q, E, w)$. Given $I \subset Q$, the *subgraph induced by the restriction of G to I*, denoted $G \cap I$, is the graph with vertex set equal to I and with edge set containing all weighted edges of G where both vertices belong to I, i.e. $(Q, E, w) \cap I = (Q \cap I, E \cap (I \times I), w|_{I \times I})$. The induced subgraph is a weighted graph with a notion of distance between vertices: given $h, k \in I$, we write $d_I(h, k) := d_{G \cap I}(h, k)$. Note that $d_I(h, k) \geq d_G(h, k)$.

We then define a *connected subset of Q* as a subset $S \subset Q$ such that $S \neq \emptyset$ and $G \cap S$ is connected, with $\mathcal{C}(Q)$ denoting the set of such subsets. We can then define partitions of Q into connected sets as follows.

Definition II.1 (Connected partitions) *Given the graph $G = (Q, E, w)$, we define a* connected $N-$partition of Q *as a collection $p = \{p_i\}_{i=1}^N$ of N subsets of Q such that*

(i) $\bigcup_{i=1}^N p_i = Q$;
(ii) $p_i \cap p_j = \emptyset$ if $i \neq j$;
(iii) $p_i \neq \emptyset$ for all $i \in \{1, \ldots, N\}$;
(iv) $p_i \in \mathcal{C}(Q)$ for all $i \in \{1, \ldots, N\}$.

Let $\mathcal{P}$ to be the set of such partitions.

Remark II.2 When $|Q| \geq N$ and G is connected, it is always possible to find a connected $N-$partition of Q. Pick N unique vertices from Q to seed the subsets, then iteratively grow each subset by adding unclaimed neighboring vertices.

For our gossip algorithms we need to introduce the notion of adjacent subsets. Two distinct connected subsets p_i, p_j are said to be *adjacent* if there are two vertices q_i, q_j belonging, respectively, to p_i and p_j such that $(q_i, q_j) \in E$. Observe that if p_i and p_j are adjacent then $p_i \cup p_j \in \mathcal{C}(Q)$. Similarly, we say that robots i and j are adjacent or are neighbors if their subsets p_i and p_j are adjacent. Accordingly we next provide the definition of an *adjacency graph*, which we will use in the analysis of our gossip algorithms.

Definition II.3 (Adjacency graph) *For $p \in \mathcal{P}$, we define the* adjacency graph *between regions of partition p as $\mathcal{G}(p) = (\{1, \ldots, N\}, \mathcal{E}(p))$, where $(i, j) \in \mathcal{E}(p)$ if p_i and p_j are adjacent.*

On Q, we define a *weight function* to be a bounded positive function $\phi : Q \to \mathbb{R}_{>0}$ which assigns a positive weight to each element of Q. Given $p_i \in \mathcal{C}(Q)$, we define the *one-center function* $\mathcal{H}_1 : p_i \to \mathbb{R}_{\geq 0}$ as

$$\mathcal{H}_1(h; p_i) = \sum_{k \in p_i} d_{p_i}(h, k) \phi(k).$$

A technical assumption is then needed to define the *generalized centroid* of a connected subset. In what follows, we assume that a *total order* relation, $<$, is defined on Q: hence, we can also denote $Q = \{1, \ldots, |Q|\}$.

Definition II.4 (Centroid) *Let Q be a totally ordered set, and $p_i \in \mathcal{C}(Q)$. We define generalized centroid of p_i as*

$$\mathrm{Cd}(p_i) := \min\{\underset{h \in p_i}{\mathrm{argmin}}\, \mathcal{H}_1(h; p_i)\}.$$

In subsequent use we will drop the word "generalized" for brevity. Note that with this definition the centroid is well-defined, and also that the centroid of a region always belongs to the region. With a slight notational abuse, we define $\mathrm{Cd} : \mathcal{P} \to Q^N$ as the map which associates to a partition the vector of the centroids of its elements.

With these notions we can introduce the *multicenter function* $\mathcal{H}_{\mathrm{multicenter}} : \mathcal{P} \times Q^N \to \mathbb{R}_{\geq 0}$ defined by

$$\mathcal{H}_{\mathrm{multicenter}}(p, c) = \sum_{i=1}^N \mathcal{H}_1(c_i; p_i).$$

Our motivation for using this cost function is to optimize the response time of the robots to a task appearing randomly in Q according relative weights ϕ. We aim to minimize $\mathcal{H}_{\mathrm{multicenter}}$ with respect to both the partitions p and the points c so as to minimize the expected distance from this random vertex to the centroid of the partition the vertex is in.

Among the ways of partitioning Q, there is one which is worth special attention. Given points $c \in Q^N$ such that if $i \neq j$, then $c_i \neq c_j$, the partition $p \in \mathcal{P}$ is said to be a *Voronoi partition of Q generated by c* if, for each p_i and all $k \in p_i$, we have $c_i \in p_i$ and $d_G(k, c_i) \leq d_G(k, c_j)$, $\forall j \neq i$.

Proposition II.5 (Properties of multicenter function) *Let $p \in \mathcal{P}$, $c \in Q^N$, and let p^* be a Voronoi partition generated by c. Then*

$$\mathcal{H}_{\mathrm{multicenter}}(p, \mathrm{Cd}(p)) \leq \mathcal{H}_{\mathrm{multicenter}}(p, c),$$
$$\mathcal{H}_{\mathrm{multicenter}}(p^*, c) \leq \mathcal{H}_{\mathrm{multicenter}}(p, c).$$

These statements motivate the following definition: a partition $p \in \mathcal{P}$ is a *centroidal Voronoi partition* if p is a Voronoi

partition generated by $\mathrm{Cd}(p)$. Based on the multicenter function, we define $\mathcal{H}_{\text{expected}} : \mathcal{P} \to \mathbb{R}_{\geq 0}$ by

$$\mathcal{H}_{\text{expected}}(p) = \mathcal{H}_{\text{multicenter}}(p, \mathrm{Cd}(p))$$
$$= \sum_{i=1}^{N} \sum_{x \in p_i} d_{p_i}(x, \mathrm{Cd}(p_i))\phi(x).$$

Observe that $\mathcal{H}_{\text{expected}}$ has the following property as an immediate consequence of Proposition II.5: given $p \in \mathcal{P}$, if p^* is a Voronoi partition generated by $\mathrm{Cd}(p)$ then

$$\mathcal{H}_{\text{expected}}(p^*) \leq \mathcal{H}_{\text{expected}}(p).$$

We are now in a position to state the goal of our algorithm: solving the optimization problem

$$\min_{p \in \mathcal{P}} \mathcal{H}_{\text{expected}}(p),$$

using only pairwise territory exchanges between agents along the edges of $\mathcal{E}(p)$. In the literature this optimization is typically studied with either centralized control or with synchronous and reliable communication between several agents, see [3] and [9]. We believe these communication requirements are unrealistic for deployed robotic networks, and thus are interested in solutions requiring only pairwise gossip communication as first explored in [13] and [14].

III. LLOYD-TYPE GOSSIP GRAPH COVERAGE

In this Section we briefly review the *discretized Lloyd-type gossip coverage algorithm* proposed in [14].

Lloyd-type Gossip Graph Coverage Algorithm

At each time $t \in \mathbb{Z}_{\geq 0}$, each agent $i \in \{1, \dots, N\}$ maintains in memory a connected subset $p_i(t)$. The collection $p(0) = \{p_1(0), \dots, p_N(0)\}$ is an arbitrary connected N−partition of Q. At each $t \in \mathbb{Z}_{\geq 0}$ a communicating pair, say $(i, j) \in \mathcal{E}(p(t))$, is selected by a deterministic or stochastic process to be determined. Assume that $i < j$. Every agent $k \notin \{i, j\}$ sets $p_k(t+1) = p_k(t)$, while i and j perform the following:

1: agent i transmits its subset $p_i(t)$ to j and vice-versa
2: both agents compute centroids $\mathrm{Cd}(p_i)$ and $\mathrm{Cd}(p_j)$, set $u = p_i \cup p_j$, and the sets

$$W_{i \to j} = \{x \in p_i : d_u(x, \mathrm{Cd}(p_j)) < d_u(x, \mathrm{Cd}(p_i))\}$$
$$W_{j \to i} = \{x \in p_j : d_u(x, \mathrm{Cd}(p_i)) < d_u(x, \mathrm{Cd}(p_j))\}$$
$$W_{i \cong j} = \{x \in p_i \cup p_j : d_u(x, \mathrm{Cd}(p_i)) = d_u(x, \mathrm{Cd}(p_j))\}$$

3: **if** $W_{i \to j} \cup W_{j \to i} = \emptyset$ **then**
4:    $p_i(t+1) := p_i(t)$ and $p_j(t+1) := p_j(t)$
5: **else**
6:    $p_i(t+1) := ((p_i \setminus W_{i \to j}) \cup W_{j \to i}) \cup W_{i \cong j},$
      $p_j(t+1) := ((p_j \setminus W_{j \to i}) \cup W_{i \to j}) \setminus (W_{i \cong j} \cap p_j)$
7: **end if**

Observe that $W_{i \to j}$ (resp. $W_{j \to i}$) contains the cells of p_i (resp. p_j) which are closer to $\mathrm{Cd}(p_j)$ (resp. $\mathrm{Cd}(p_i)$), whereas $W_{i \cong j}$ represents the set of the tied cells. In other words,

when two robots exchange territory using this Lloyd-type algorithm, their updated regions $\{p_i(t+1), p_j(t+1)\}$ are the Voronoi partition of the set $p_i(t) \cup p_j(t)$ generated by the centroids $\mathrm{Cd}(p_i)$ and $\mathrm{Cd}(p_j)$, with all tied cells assigned to the agent with the lower index.

Now, for any pair $(i, j) \in \{1, \dots, N\}^2$, $i \neq j$, we define the map $T_{ij} : \mathcal{P} \rightrightarrows \mathcal{P}$ as

$$T_{ij}(p) = \begin{cases} p, & \text{if } (i,j) \notin \mathcal{E}(p) \text{ or } W_{i \to j} \cup W_{j \to i} = \emptyset \\ \{p_1, \dots, \widehat{p}_i, \dots, \widehat{p}_j, \dots, p_N\}, & \text{otherwise,} \end{cases}$$

where $\widehat{p}_i = ((p_i \setminus W_{i \to j}) \cup W_{j \to i}) \cup W_{i \cong j}$, and $\widehat{p}_j = ((p_j \setminus W_{j \to i}) \cup W_{i \to j}) \setminus (W_{j \cong i} \cap p_j)$.

The dynamical system on the space of partitions is therefore described by, for $t \in \mathbb{Z}_{\geq 0}$,

$$p(t+1) = T_{ij}(p(t)), \quad \text{for some } (i,j) \in \mathcal{E}(p(t)), \quad (1)$$

together with a rule for which edge (i, j) is selected at each time. We also define the set-valued map $T : \mathcal{P} \rightrightarrows \mathcal{P}$ by

$$T(p) = \{T_{ij}(p) \mid (i,j) \in \mathcal{E}(p)\}. \quad (2)$$

In Theorem (III.2) we summarize the convergence properties of the Lloyd-type discretized gossip coverage algorithm. To do so, we need the following definition.

Definition III.1 (Uniform and random persistency) *Let X be a finite set.*
(i) *A map $\sigma : \mathbb{Z}_{\geq 0} \to X$ is* uniformly persistent *if there exists a duration $\Delta \in \mathbb{N}$ such that, for each $x \in X$, there exists an increasing sequence of times $\{t_k\}_{k \in \mathbb{Z}_{\geq 0}} \subset \mathbb{Z}_{\geq 0}$ satisfying $t_{k+1} - t_k \leq \Delta$ and $\sigma(t_k) = x$ for all $k \in \mathbb{Z}_{\geq 0}$.*
(ii) *A stochastic process $\sigma : \mathbb{Z}_{\geq 0} \to X$ is* randomly persistent *if there exists a probability $p \in \,]0, 1[$ such that, for each $x \in X$ and for each $t \in \mathbb{Z}_{\geq 0}$*

$$\mathbb{P}\big[\sigma(t+1) = x \,|\, \sigma(t), \dots, \sigma(1)\big] \geq p.$$

Theorem III.2 (Convergence under persistent gossip) *Consider the Lloyd-type discretized gossip coverage algorithm T and the evolutions $p : \mathbb{Z}_{\geq 0} \to \mathcal{P}$ defined by*

$$p(t+1) = T_{\sigma(t)}(p(t)), \quad \text{for } t \in \mathbb{Z}_{\geq 0},$$

where $\sigma : \mathbb{Z}_{\geq 0} \to \{(i, j) \in \{1, \dots, N\}^2 \mid i \neq j\}$ is either a deterministic map or a stochastic process. Then the following statements hold:
(i) *if σ is a uniformly persistent map, then any evolution p converges in a finite number of steps to a centroidal Voronoi partition; and*
(ii) *if σ is a randomly persistent stochastic process, then any evolution p converges almost surely in a finite number of steps to a centroidal Voronoi partition.*

Convergence to a centroidal Voronoi partition puts this pairwise optimization algorithm in the same class with centralized Lloyd methods. However, there can be a great number of possible centroidal Voronoi partitions for a given discretized environment and, in our experience, the algorithm too frequently converges to a suboptimal solution. This issue motivates the developments in the following Section.

IV. Pairwise-optimal gossip coverage algorithm

In this section we present a novel gossip graph coverage algorithm which improves on the performance achievable by the Lloyd-type algorithm reviewed in the previous Section. First we introduce the following notational definition.

Given $p \in \mathcal{P}$ and distinct components p_i and p_j, let $P_{p_i \cup p_j}$ denote the set of all distinct pairs of vertices in $p_i \cup p_j$. We assume that the elements in $P_{p_i \cup p_j}$ are sorted lexicographically. In formal terms, if $p_i \cup p_j = \{h_1, \ldots, h_{|p_i \cup p_j|}\}$ where $h_s < h_{s+1}$ for $s = \{1, \ldots, |p_i \cup p_j| - 1\}$, then

$$P_{p_i \cup p_j} = \big\{(h_1, h_2), (h_1, h_3), \ldots, (h_1, h_{|p_i \cup p_j|}), (h_2, h_3), \ldots$$
$$\ldots, (h_{|p_i \cup p_j| - 1}, h_{|p_i \cup p_j|})\big\}.$$

Notice that $|P_{p_i \cup p_j}| = \frac{|p_i \cup p_j| \, (|p_i \cup p_j| - 1)}{2}$.

The *pairwise-optimal discretized gossip coverage algorithm* is defined as follows.

Gossip Optimal Graph Coverage Algorithm

At each time $t \in \mathbb{Z}_{\geq 0}$, each agent $i \in \{1, \ldots, N\}$ maintains in memory a connected subset $p_i(t)$. The collection $p(0) = \{p_1(0), \ldots, p_N(0)\}$ is an arbitrary connected $N-$partition of Q. At each $t \in \mathbb{Z}_{\geq 0}$ a communicating pair, say $(i, j) \in \mathcal{E}(p(t))$, is selected by a deterministic or stochastic process to be determined. Assume that $i < j$. Every agent $k \notin \{i, j\}$ sets $p_k(t+1) = p_k(t)$, while i and j perform the following:

1: agent i transmits its subset $p_i(t)$ to j and vice-versa
2: both agents compute the set $u = p_i \cup p_j$, the function $\mathcal{D} : P_u \to \mathbb{R}_{>0}$

$$\mathcal{D}((a, b)) = \sum_{h \in u} \min \{d_u(h, a), d_u(h, b)\},$$

and find the set $\mathcal{S}_\mathcal{D} = \mathrm{argmin}_{(h_s, h_r) \in P_u} \{\mathcal{D}(h_s, h_r)\}$
3: assuming $\mathcal{S}_\mathcal{D}$ is lexicographically ordered, both agents pick the first pair in $\mathcal{S}_\mathcal{D}$, say (a^*, b^*)
4: both agents compute the sets

$$W_{a^*} = \{x \in u : d_u(x, a^*) \leq d_u(x, b^*)\}$$
$$W_{b^*} = \{x \in u : d_u(x, b^*) < d_u(x, a^*)\}$$

5: **if** $\mathcal{H}_1(\mathrm{Cd}(W_{a^*}); W_{a^*}) + \mathcal{H}_1(\mathrm{Cd}(W_{b^*}); W_{b^*}) < \mathcal{H}_1(\mathrm{Cd}(p_i(t)); p_i(t)) + \mathcal{H}_1(\mathrm{Cd}(p_j(t)); p_j(t))$ **then**
6:     $p_i(t+1) = W_{a^*}, \quad p_j(t+1) = W_{b^*}$
7: **else**
8:     $p_i(t+1) = p_i(t), \quad p_j(t+1) = p_j(t)$
9: **end if**

Remark IV.1 The lexicographic rule for picking (a^*, b^*) used here makes the dynamical system represented by the algorithm deterministic and well-defined. In practice any pair in $\mathcal{S}_\mathcal{D}$ can be chosen.

The following Proposition characterizes a desirable property of the regions W_{a^*} and W_{b^*}.

Proposition IV.2 *The regions W_{a^*} and W_{b^*} as defined in Gossip Optimal Graph Coverage Algorithm, are connected.*

Thanks to Proposition IV.2 we have that, for any pair $(i, j) \in \{1, \ldots, N\}^2$, $i \neq j$, we can define the map $T_{ij} : \mathcal{P} \to \mathcal{P}$ by

$$T_{ij}(p) = \begin{cases} p, & \text{if} \quad (i, j) \notin \mathcal{E}(p) \\ \{p_1, \ldots, \widehat{p}_i, \ldots, \widehat{p}_j, \ldots, p_N\}, & \text{otherwise,} \end{cases}$$

where $\widehat{p}_i = W_{a^*}$ and $\widehat{p}_j = W_{b^*}$. Therefore, the dynamical system on the space of partitions can again be described in compact form by (1) and (2) but with this new T_{ij}.

To establish the convergence properties of the *pairwise-optimal discretized gossip coverage algorithm* we need the following definition.

Definition IV.3 *A partition $p \in \mathcal{P}$ is said to be a* pairwise-optimal partition *if, for every $(i, j) \in \mathcal{E}(p)$*

$$\mathcal{H}_1(\mathrm{Cd}(p_i); p_i) + \mathcal{H}_1(\mathrm{Cd}(p_j); p_j) =$$
$$\min_{a, b \in p_i \cup p_j} \left\{ \sum_{k \in p_i \cup p_j} \min \{d_{p_i \cup p_j}(a, k), d_{p_i \cup p_j}(b, k)\} \right\}$$

The following Proposition characterizes the set of the pairwise-optimal partitions.

Proposition IV.4 *Let $p \in \mathcal{P}$ be a pairwise-optimal partition. Then p is also a centroidal Voronoi partition.*

We are now ready to state the main result of this paper.

Theorem IV.5 (Convergence under persistent gossip)
Consider the pairwise-optimal discretized gossip coverage algorithm T and the evolutions $p : \mathbb{Z}_{\geq 0} \to \mathcal{P}$ defined by

$$p(t+1) = T_{\sigma(t)}(p(t)), \quad \text{for } t \in \mathbb{Z}_{\geq 0},$$

where $\sigma : \mathbb{Z}_{\geq 0} \to \{(i, j) \in \{1, \ldots, N\}^2 \mid i \neq j\}$ is either a deterministic map or a stochastic process. Then the following statements hold:

(i) *if σ is a uniformly persistent map, then any evolution p converges in a finite number of steps to a pairwise-optimal partition; and*

(ii) *if σ is a randomly persistent stochastic process, then any evolution p converges almost surely in a finite number of steps to a pairwise-optimal partition.*

Proofs for these statements are provided in [15].

In contrast to the Lloyd-type algorithm in Section III, this new pairwise-optimal algorithm converges to an element in the set of pairwise-optimal partitions. This result is a useful improvement as the set of pairwise-optimal partitions is a subset of the set of centroidal Voronoi partitions, so that the new algorithm can avoid many of the local minima in which Lloyd-type methods get stuck. We will demonstrate this improvement in Section VI.
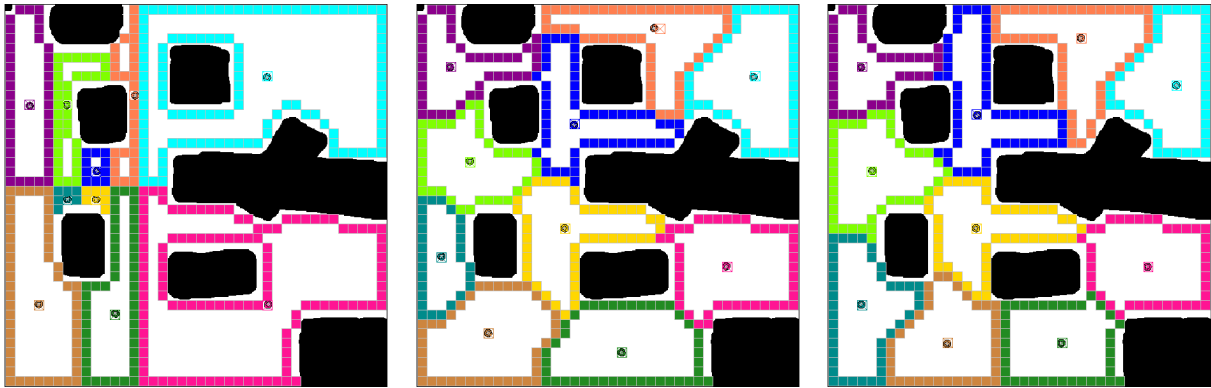
Fig. 1. At left, ten robots are positioned at the centroids of their initial partitions in a non-convex environment. The boundary of each agent's partition is drawn in a different color. The middle shows the final solution of a centralized Lloyd optimization with a total cost of $624.1m$. At right is a final solution for the pairwise-optimal discretized gossip coverage algorithm starting from the final centralized Lloyd solution, with an improved cost of $610.3m$.

V. SCALABILITY PROPERTIES

In this Section we discuss how to compute the optimal new centroids $(a^*, b^*) \in P_{p_i \cup p_j}$ for the Gossip Optimal Graph Coverage Algorithm, as well as a sampling method to greatly reduce the computational complexity. Determining (a^*, b^*) requires an exhaustive search over all possible pairs of vertices for a pair with the lowest cost to cover $p_i \cup p_j$. Using Johnson's all pairs shortest paths algorithm as a first step to compute a pairwise distance matrix for $p_i \cup p_j$, (a^*, b^*) can be found in $\mathcal{O}(|p_i|^3)$ with a memory requirement of $\mathcal{O}(|p_i|^2)$ (as $p_i \cup p_j$ is of order $|p_i|$).

For mobile robots with limited onboard resources, these requirements may be too steep. In such circumstances, we propose instead to sample pairs of potential new centroids when applying the algorithm. The first sample pair would be the agents' previous centroids, with the rest drawn at random from the set $P_{p_i \cup p_j}$. For m sample pairs, this approach requires $2m$ calls to Dijkstra's one-to-all shortest paths algorithm to find the pair with the lowest cost, for a total time complexity of $\mathcal{O}(m|p_i| \log(|p_i|))$ and a memory requirement of $\mathcal{O}(|p_i|)$. It is also worth noting that if all edge weights in G are equal (i.e., for a regular grid discretization), then a breadth-first-search approach can replace Dijkstra and the time complexity drops to $\mathcal{O}(m|p_i|)$. This sample-based approach greatly reduces the time and memory requirements, and will still converge to a pairwise-optimal partition.

VI. SIMULATION RESULTS

To demonstrate the utility of the proposed gossip coverage algorithm, we implemented it and other coverage algorithms using the open-source Player/Stage robot software system [16] and the Boost Graph Library (BGL) [17]. All results presented here were generated using Player 2.1.1, Stage 2.1.0, and BGL 1.34.1. A non-convex environment was specified with a bitmap and overlaid with a $0.1m$ resolution occupancy grid, producing a lattice-like graph with all edge weights equal to $0.1m$. To compute distances on graphs with uniform edge weights we extended the BGL implementation of breadth-first-search with a distance recorder event visitor.
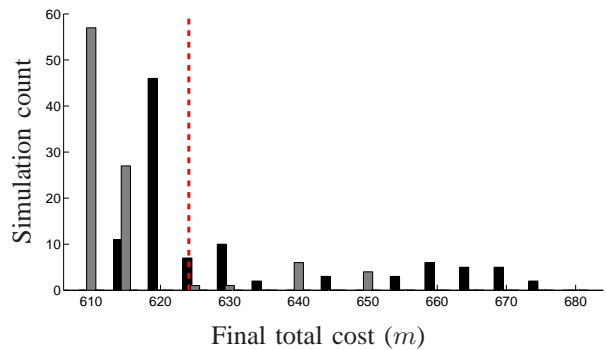

Fig. 2. Histograms of final total costs for 100 simulations of pairwise-optimal (gray) and Lloyd-type (black) gossip coverage. All runs started from the initial partition in Fig. 1 with different random sequences of agent pairs. The dashed red line shows the final cost using a centralized Lloyd algorithm.

Figure 1 provides one example of how convergence to the set of pairwise-optimal partitions represents an improvement over Lloyd-type methods. On the left is the non-convex environment used for all results in this Section, as well as an initial partition for 10 robots. The middle panel shows the result of a centralized Lloyd optimization where all agents iteratively update their centroid and partition based on the centroids of all other agents. The final equilibrium for this centralized method is a centroidal Voronoi partition with a total cost of $624.1m$. However, this solution is not a pairwise-optimal partition. We started the pairwise-optimal discretized gossip coverage algorithm using this partition as the initial condition. After 70 pairwise exchanges between random agent pairs, the algorithm reached the pairwise-optimal partition at right which has a lower cost of 610.3. The best known 10-partition of this map has cost 610.0.

Figure 2 compares 100 simulations of the Lloyd-type and pairwise-optimal discretized gossip coverage algorithms. We started each simulation from the initial partition in Fig. 1 and used different random sequences of agent pairs for each case. The new pairwise-optimal algorithm shows marked improvement over the Lloyd-type gossip algorithm, achieving a total cost within 2% of the best known cost in 85% of trials, while also avoiding the worst local minima.
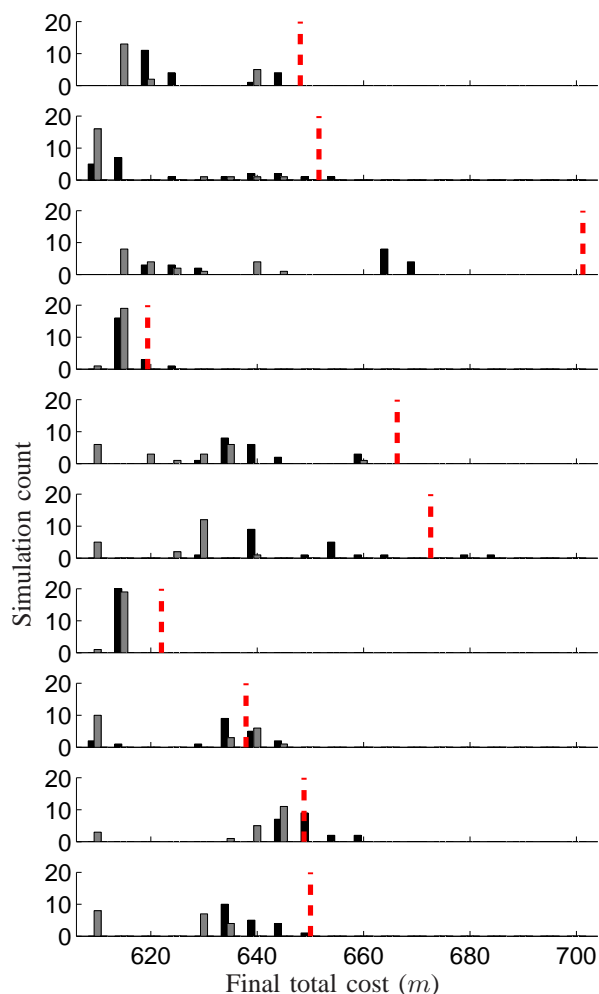
Fig. 3. Histograms of final costs for simulations starting from 10 random initial conditions for the environment in Fig. 1. The pairwise-optimal (gray) and Lloyd-type (black) gossip coverage algorithms each ran 20 times per initial condition using different random sequences of agent pairs. The dashed red line shows the final cost using a centralized Lloyd algorithm.

Figure 3 compares final cost histograms for 10 random initial conditions for the same environment. Each initial condition was created by selecting unique starting locations for the agents uniformly at random, and using these locations to generate a Voronoi partition. The histograms compare 20 simulations of Lloyd-type and the new pairwise-optimal gossip coverage starting from the same initial partition but with different random orders of exchanges. The centralized Lloyd final cost for each initial condition is also shown. The new pairwise-optimal method outperforms both Lloyd-type methods for all 10 tests, although the Lloyd-type gossip method is close in two of the trials. These results illustrate that convergence to a pairwise-optimal partition represents a significant performance enhancement over classic Lloyd methods. Interestingly, the Lloyd-type gossip algorithm also substantially outperforms the centralized version in 8 of the trials. We speculate this is due to the gossip method taking trajectories through the space of connected $N-$partitions which are not possible for the centralized approach.

VII. CONCLUSIONS

We have presented a novel distributed coverage control algorithm which requires only pairwise communication between agents. The classic Lloyd approach to optimizing quantizer placement involves iteration of separate centering and Voronoi partitioning steps. In the graph coverage domain, the separation of centroid and partition optimization seems unnecessary computationally for gossip algorithms. More importantly, we have shown that improved performance can be achieved without this separation. Our new pairwise-optimal discretized gossip coverage algorithm provably converges to a subset of the centroidal Voronoi partitions which we labeled the set of pairwise-optimal partitions. Through numerical comparisons we demonstrated that this new subset of solutions avoids many of the local minima in which both gossip and centralized Lloyd-type algorithms can get stuck.

REFERENCES

[1] E. Fiorelli, N. E. Leonard, P. Bhatta, D. A. Paley, R. Bachmayer, and D. M. Fratantoni, "Multi-AUV control and adaptive sampling in Monterey Bay," *IEEE Journal of Oceanic Engineering*, vol. 31, no. 4, pp. 935–948, 2006.

[2] P. R. Wurman, R. D'Andrea, and M. Mountz, "Coordinating hundreds of cooperative, autonomous vehicles in warehouses," *AI Magazine*, vol. 29, no. 1, pp. 9–20, 2008.

[3] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*. Applied Mathematics Series, Princeton University Press, 2009. Available at http://www.coordinationbook.info.

[4] S. P. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982. Presented as Bell Laboratory Technical Memorandum at a 1957 Institute for Mathematical Statistics meeting.

[5] C. Gao, J. Cortés, and F. Bullo, "Notes on averaging over acyclic digraphs and discrete coverage control," *Automatica*, vol. 44, no. 8, pp. 2120–2127, 2008.

[6] V. V. Vazirani, *Approximation Algorithms*. Springer, 2001.

[7] M. Zhong and C. G. Cassandras, "Distributed coverage control in sensor network environments with polygonal obstacles," in *IFAC World Congress*, (Seoul, Korea), pp. 4162–4167, July 2008.

[8] C. H. Caicedo-Nùñez and M. Žefran, "Performing coverage on nonconvex domains," in *IEEE Conf. on Control Applications*, (San Antonio, TX), pp. 1019–1024, Sept. 2008.

[9] L. C. A. Pimenta, V. Kumar, R. C. Mesquita, and G. A. S. Pereira, "Sensing and coverage for a network of heterogeneous robots," in *IEEE Conf. on Decision and Control*, (Cancún, México), pp. 3947–3952, Dec. 2008.

[10] O. Baron, O. Berman, D. Krass, and Q. Wang, "The equitable location problem on the plane," *European Journal of Operational Research*, vol. 183, no. 2, pp. 578–590, 2007.

[11] S. Yun, M. Schwager, and D. Rus, "Coordinating construction of truss structures using distributed equal-mass partitioning," in *International Symposium on Robotics Research*, (Lucerne, Switzerland), Aug. 2009.

[12] M. Schwager, D. Rus, and J. J. Slotine, "Decentralized, adaptive coverage control for networked robots," *International Journal of Robotics Research*, vol. 28, no. 3, pp. 357–375, 2009.

[13] F. Bullo, R. Carli, and P. Frasca, "Gossip coverage control for robotic networks: Dynamical systems on the the space of partitions," Dec. 2009. Available at http://arxiv.org/abs/0903.3642.

[14] J. W. Durham, R. Carli, P. Frasca, and F. Bullo, "Discrete partitioning and coverage control with gossip communication," in *ASME Dynamic Systems and Control Conference*, (Hollywood, CA), Oct. 2009.

[15] J. W. Durham, R. Carli, and F. Bullo, "Pairwise optimal discrete coverage control for gossiping robots," Aug. 2010. Available at http://arxiv.org/abs/1008.4941.

[16] B. Gerkey, R. T. Vaughan, and A. Howard, "The Player/Stage Project." http://playerstage.sourceforge.net, June 2008. Version 2.11.

[17] J. G. Siek, L.-Q. Lee, and A. Lumsdaine, "Boost Graph Library." http://www.boost.org, July 2007. Version 1.34.1.