

Discrete Partitioning and Coverage Control for Gossiping Robots

Joseph W. Durham Ruggero Carli Paolo Frasca Francesco Bullo

Abstract—We propose distributed algorithms to automatically deploy a team of mobile robots to partition and provide coverage of a non-convex environment. To handle arbitrary non-convex environments, we represent them as graphs. Our partitioning and coverage algorithm requires only short-range, unreliable pairwise “gossip” communication. The algorithm has two components: (1) a motion protocol to ensure that neighboring robots communicate at least sporadically, and (2) a pairwise partitioning rule to update territory ownership when two robots communicate. By studying an appropriate dynamical system on the space of partitions of the graph vertices, we prove that territory ownership converges to a pairwise-optimal partition in finite time. This new equilibrium set represents improved performance over common Lloyd-type algorithms. Additionally, we detail how our algorithm scales well for large teams in large environments and how the computation can run in anytime with limited resources. Finally, we report on large-scale simulations in complex environments and hardware experiments using the Player/Stage robot control system.

I. INTRODUCTION

Coordinated networks of mobile robots are already in use for environmental monitoring and warehouse logistics. In the near future, autonomous robotic teams will revolutionize transportation of passengers and goods, search and rescue operations, and other applications. These tasks share a common feature: the robots are asked to provide service over a space. One question which arises is: when a group of robots is waiting for a task request to come in, how can they best position themselves to be ready to respond?

The distributed *environment partitioning problem* for robotic networks consists of designing individual control and communication laws such that the team divides a large space into regions. Typically, partitioning is done so as to optimize a cost function which measures the quality of service provided over all of the regions. *Coverage control* additionally optimizes the positioning of robots inside a region as shown in Fig. 1.

This paper describes a distributed partitioning and coverage control algorithm for a network of robots to minimize the

This work was supported in part by ARO MURI Award W911NF-05-1-0219, NSF grants IIS-0904501 and CPS-1035917, and MIUR grant PRIN-20087W5P2. Preliminary and incomplete versions of this work appeared in the Proceedings of the 2009 ASME Dynamic Systems and Control Conference, Hollywood, California, USA, and in the Proceedings of the 2010 IEEE Conference on Decision and Control, Atlanta, Georgia, USA.

Joseph W. Durham and Francesco Bullo are with the Department of Mechanical Engineering, University of California, Santa Barbara, CA, 93106 (joe, bullo)@engineering.ucsb.edu

Ruggero Carli is with the Department of Information Engineering, University of Padova, Via Gradenigo 6/a, 35131 Padova, Italy carlirug@dei.unipd.it

Paolo Frasca is with the Dipartimento di Matematica, Politecnico di Torino, corso Duca degli Abruzzi 24, 10129 Torino, Italy paolo.frasca@polito.it

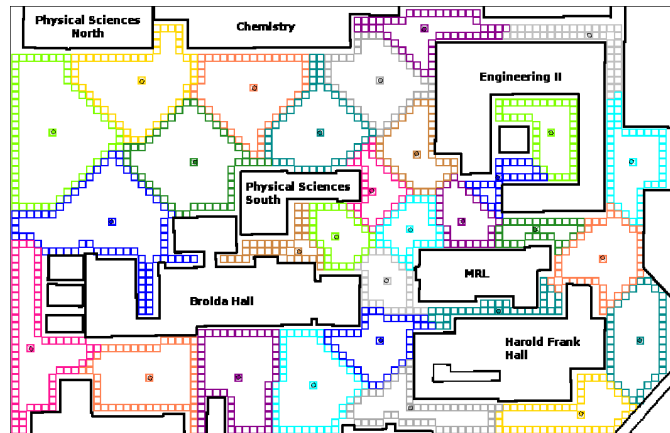


Fig. 1. Example of a team of robots providing efficient coverage of a non-convex environment, as measured by an appropriate multicenter cost function.

expected distance between the closest robot and spatially distributed events which will appear at discrete points in a non-convex environment. Optimality is defined with reference to a relevant “multicenter” cost function. As with all multirobot coordination applications, the challenge comes from reducing the communication requirements: the proposed algorithm requires only short-range “gossip” communication, i.e., asynchronous and unreliable communication between nearby robots.

Literature Review

Territory partitioning and coverage control have applications in many fields. In cyber-physical systems, applications include automated environmental monitoring [1], fetching and delivery [2], construction [3], and other vehicle routing scenarios [4]. More generally, coverage of discrete sets is also closely related to the literature on data clustering and k -means [5], as well as the facility location or k -center problem [6]. Partitioning of graphs is its own field of research, see [7] for a survey. Territory partitioning through local interactions is also studied for animal groups, see for example [8].

A broad discussion of algorithms for partitioning and coverage control in robotic networks is presented in [9] which builds on the classic work of Lloyd [10] on optimal quantizer selection through “centering and partitioning.” The Lloyd approach was first adapted for distributed coverage control in [11]. Since this beginning, similar algorithms have been applied to non-convex environments [12], [13], unknown density functions [14], [15], equitable partitioning [16], and construction of truss-like objects [3]. There are also multi-agent partitioning algorithms built on market principles or

auctions, see [17] for a survey.

While Lloyd iterative optimization algorithms are popular and work well in simulation, they require synchronous and reliable communication among neighboring robots. As robots with adjacent regions may be arbitrarily far apart, these communication requirements are burdensome and unrealistic for deployed robotic networks. In response to this issue, in [18] the authors have shown how a group of robotic agents can optimize the partition of a convex bounded set using a Lloyd algorithm with gossip communication. A Lloyd algorithm with gossip communication has also been applied to optimizing partitions of non-convex environments in [19], the key idea being to transform the coverage problem in Euclidean space into a coverage problem on a graph with geodesic distances.

Distributed Lloyd methods are built around separate partitioning and centering steps, and they are attractive because there are known ways to characterize their equilibrium sets (the so-called centroidal Voronoi partitions) and prove convergence. Unfortunately, even for very simple environments (both continuous and discrete) the set of centroidal Voronoi partitions may contain several sub-optimal configurations. We are thus interested in studying (discrete) gossip coverage algorithms for two reasons: (1) they apply to more realistic robot network models featuring very limited communication in large non-convex environments, and (2) they are more flexible than typical Lloyd algorithms meaning they can avoid poor suboptimal configurations and improve performance.

Statement of Contributions

There are three main contributions in this paper. First, we present a discrete partitioning and coverage optimization algorithm for mobile robots with unreliable, asynchronous, and short-range communication. Our algorithm has two components: a *motion protocol* which drives the robots to meet their neighbors, and a *pairwise partitioning rule* to update territories when two robots meet. The partitioning rule optimizes coverage of a set of points connected by edges to form a graph. The flexibility of graphs allows the algorithm to operate in non-convex, non-polygonal environments with holes. Our graph partition optimization approach can also be applied to non-planar problems, existing transportation or logistics networks, or more general data sets.

Second, we provide an analysis of both the convergence properties and computational requirements of the algorithm. By studying a dynamical system of partitions of the graph's vertices, we prove that almost surely the algorithm converges to a pairwise-optimal partition in finite time. The set of pairwise-optimal partitions is shown to be a proper subset of the well-studied set of centroidal Voronoi partitions. We further describe how our pairwise partitioning rule can be implemented to run in anytime and how the computational requirements of the algorithm can scale up for large domains and large teams.

Third, we detail experimental results from our implementation of the algorithm in the Player/Stage robot control system. We present a simulation of 30 robots providing coverage of a portion of a college campus to demonstrate that our algorithm

can handle large robot teams, and a hardware-in-the-loop experiment conducted in our lab which incorporates sensor noise and uncertainty in robot position. Through numerical analysis we also show how our new approach to partitioning represents a significant performance improvement over both common Lloyd-type methods and the recent results in [18].

The present work differs from the gossip Lloyd method [18] in three respects. First, while [18] focuses on territory partitioning in a convex continuous domain, here we operate on a graph which allows our approach to consider geodesic distances, work in non-convex environments, and maintain connected territories. Second, instead of a pairwise Lloyd-like update, we use an iterative optimal two-partitioning approach which yields better final solutions. Third, we also present a motion protocol to produce the sporadic pairwise communications required for our gossip algorithm and characterize the computational complexity of our proposal. Preliminary versions of this paper appeared in [19] and [20]. Compared to these, the new content here includes: (1) a motion protocol; (2) a simplified and improved pairwise partitioning rule; (3) proofs of the convergence results; and (4) a description of our implementation and a hardware-in-the-loop experiment.

Paper Structure and Notation

In Section II we review and adapt coverage and geometric concepts (e.g., centroids, Voronoi partitions) to a discrete environment like a graph. We formally describe our robot network model and the discrete partitioning problem in Section III, and then state our coverage algorithm and its properties. Section IV contains proofs of the main convergence results. In Section V we detail our implementation of the algorithm and present experiments and comparative analysis. Some conclusions are given in Section VI.

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 A , $|A|$ denotes the number of elements in A . Given sets A, B , their difference is $A \setminus B = \{a \in A \mid a \notin B\}$. A set-valued map, denoted by $T : A \rightrightarrows B$, associates to an element of A a subset of B .

II. PRELIMINARIES

We are given a team of N robots tasked with providing coverage of a finite set of points in a non-convex and non-polygonal environment. In this Section we translate concepts used in coverage of continuous environments to graphs.

A. Non-convex Environment as a Graph

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

Remark 2.1 (Discretization of an Environment): For the examples in this paper we will use a coarse *occupancy grid*

map as a representation of a continuous environment. In an occupancy grid [21], each grid cell is either free space or an obstacle (occupied). To form a weighted graph, each free cell becomes a vertex and free cells are connected with edges if they border each other in the grid. Edge weights are the distances between the centers of the cells, i.e., the grid resolution. There are many other methods to discretize a space, including triangularization and other approaches from computational geometry [22], which could also be used.

In any weighted graph $G(Q)$ 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 of vertices 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 in G is finite. By convention, $d_G(h, k) = 0$ if $h = k$. Note that $d_G(h, k) = d_G(k, h)$, for any $h, k \in Q$.

B. Partitions of Graphs

We will be partitioning Q into N connected subsets or regions which will each be covered by an individual robot. To do so we need to define distances on induced subgraphs of $G(Q)$. Given $I \subset Q$, the *subgraph induced by the restriction of G to I* , denoted by $G \cap I$, is the graph with vertex set equal to I and edge set containing all weighted edges of G where both vertices belong to I . In other words, we set $(Q, E, w) \cap I = (Q \cap I, E \cap (I \times I), w|_{I \times I})$. 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 define a *connected subset* of Q as a subset $A \subset Q$ such that $A \neq \emptyset$ and $G \cap A$ is connected. We can then partition Q into connected subsets as follows.

Definition 2.2 (Connected Partitions): Given the graph $G(Q) = (Q, E, w)$, 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 $i \in \{1, \dots, N\}$; and
- (iv) P_i is connected for all $i \in \{1, \dots, N\}$.

Let $\text{Part}_N(Q)$ to be the set of connected N -partitions of Q .

Property (ii) implies that each element of Q belongs to just one P_i , i.e., each location in the environment is covered by just one robot. Notice that each $P_i \in P$ induces a connected subgraph in $G(Q)$. In subsequent references to P_i we will often mean $G \cap P_i$, and in fact we refer to $P_i(t)$ as the *dominance subgraph* or *region* of the i -th robot at time t .

Among the ways of partitioning Q , there are some which are worth special attention. Given a vector of distinct points $c \in Q^N$, the partition $P \in \text{Part}_N(Q)$ 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 $d_G(k, c_i) \leq d_G(k, c_j), \forall j \neq i$. Note that the Voronoi partition generated by c is not unique since how to apportion tied vertices is unspecified.

C. Adjacency of Partitions

For our gossip algorithms we need to introduce the notion of adjacent subgraphs. Two distinct connected subgraphs 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 $P_i \cup P_j$ is connected. Similarly, we say that robots i and j are adjacent or are neighbors if their subgraphs P_i and P_j are adjacent. Accordingly, we introduce the following useful notion.

Definition 2.3 (Adjacency Graph): For $P \in \text{Part}_N(Q)$, we define the *adjacency graph* between regions of partition P as $\mathcal{G}(P) = (\{1, \dots, N\}, \mathcal{E}(P))$, where $(i, j) \in \mathcal{E}(P)$ if P_i and P_j are adjacent.

Note that $\mathcal{G}(P)$ is always connected since $G(Q)$ is.

D. Cost Functions

We define three coverage cost functions for graphs: \mathcal{H}_{one} , $\mathcal{H}_{\text{multicenter}}$, and $\mathcal{H}_{\text{expected}}$. Let the *weight function* $\phi : Q \rightarrow \mathbb{R}_{>0}$ assign a relative weight to each element of Q . The *one-center function* \mathcal{H}_{one} gives the cost for a robot to cover a connected subset $A \subset Q$ from a vertex $h \in A$ with relative prioritization set by ϕ :

$$\mathcal{H}_{\text{one}}(h; A) = \sum_{k \in A} d_A(h, k) \phi(k).$$

A technical assumption is needed to solve the problem of minimizing $\mathcal{H}_{\text{one}}(\cdot, A)$: we assume from now on that a *total order* relation, $<$, is defined on Q , i.e., that $Q = \{1, \dots, |Q|\}$. With this assumption we can deterministically pick a vertex in A which minimizes \mathcal{H}_{one} as follows.

Definition 2.4 (Centroid): Let Q be a totally ordered set, and let $A \subset Q$. We define the set of generalized centroids of A as the set of vertices in A which minimize \mathcal{H}_{one} , i.e.,

$$C(A) := \underset{h \in A}{\text{argmin}} \mathcal{H}_{\text{one}}(h; A).$$

Further, we define the map Cd as $\text{Cd}(A) := \min\{c \in C(A)\}$. We call $\text{Cd}(A)$ the *generalized centroid* of A .

In subsequent use we 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 $\text{Cd} : \text{Part}_N(Q) \rightarrow Q^N$ as the map which associates to a partition the vector of the centroids of its elements.

We define the *multicenter function* $\mathcal{H}_{\text{multicenter}}$ to measure the cost for N robots to cover a connected N -partition P from the vertex set $c \in Q^N$:

$$\mathcal{H}_{\text{multicenter}}(c, P) = \frac{1}{\sum_{k \in Q} \phi(k)} \sum_{i=1}^N \mathcal{H}_{\text{one}}(c_i; P_i).$$

We aim to minimize the performance function $\mathcal{H}_{\text{multicenter}}$ with respect to both the vertices c and the partition P .

We can now state the coverage cost function we will be concerned with for the rest of this paper. Let $\mathcal{H}_{\text{expected}} : \text{Part}_N(Q) \rightarrow \mathbb{R}_{\geq 0}$ be defined by

$$\mathcal{H}_{\text{expected}}(P) = \mathcal{H}_{\text{multicenter}}(\text{Cd}(P), P).$$

In the motivational scenario we are considering, each robot will periodically be asked to perform a task somewhere in its region with tasks appearing according to distribution ϕ . When idle, the robots would position themselves at the centroid of their region. By partitioning G so as to minimize $\mathcal{H}_{\text{expected}}$, the robot team would minimize the expected distance between a task and the robot which will service it.

E. Optimal Partitions

We introduce two notions of optimal partitions: centroidal Voronoi and pairwise-optimal. Our discussion starts with the following simple result about the multicenter cost function.

Proposition 2.5 (Properties of Multicenter Function): Let $P \in \text{Part}_N(Q)$ and $c \in Q^N$. If P' is a Voronoi partition generated by c and $c' \in Q^N$ is such that $c'_i \in C(P_i) \forall i$, then

$$\begin{aligned} \mathcal{H}_{\text{multicenter}}(c, P') &\leq \mathcal{H}_{\text{multicenter}}(c, P), \text{ and} \\ \mathcal{H}_{\text{multicenter}}(c', P) &\leq \mathcal{H}_{\text{multicenter}}(c, P). \end{aligned}$$

The second inequality is strict if any $c_i \notin C(P_i)$.

Proposition 2.5 implies the following necessary condition: if (c, P) minimizes $\mathcal{H}_{\text{multicenter}}$, then $c_i \in C(P_i) \forall i$ and P must be a Voronoi partition generated by c . Thus, $\mathcal{H}_{\text{expected}}$ has the following property as an immediate consequence of Proposition 2.5: given $P \in \text{Part}_N(Q)$, if P^* is a Voronoi partition generated by $\text{Cd}(P)$ then

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

This fact motivates the following definition.

Definition 2.6 (Centroidal Voronoi Partition): $P \in \text{Part}_N(Q)$ is a *centroidal Voronoi partition* of Q if there exists a $c \in Q^N$ such that P is a Voronoi partition generated by c and $c_i \in C(P_i) \forall i$.

The set of *pairwise-optimal partitions* provides an alternative definition for the optimality of a partition: a partition is pairwise-optimal if, for every pair of adjacent regions, one can not find a better two-partition of the union of the two regions. This condition is formally stated as follows.

Definition 2.7 (Pairwise-optimal Partition): $P \in \text{Part}_N(Q)$ is a *pairwise-optimal partition* if for every $(i, j) \in \mathcal{E}(P)$,

$$\begin{aligned} &\mathcal{H}_{\text{one}}(\text{Cd}(P_i); P_i) + \mathcal{H}_{\text{one}}(\text{Cd}(P_j); P_j) = \\ &\min_{a, b \in P_i \cup P_j} \left\{ \sum_{k \in P_i \cup P_j} \min \{ d_{P_i \cup P_j}(a, k), d_{P_i \cup P_j}(b, k) \} \phi(k) \right\}. \end{aligned}$$

The following Proposition states that the set pairwise-optimal partitions is in fact a subset of the set of centroidal Voronoi partitions. The proof is involved and is deferred to Appendix C. See Fig. 2 for an example which demonstrates that the inclusion is strict.

Proposition 2.8 (Pairwise-optimal Implies Voronoi): Let $P \in \text{Part}_N(Q)$ be a *pairwise-optimal partition*. Then P is also a *centroidal Voronoi partition*.

For a given environment Q , a pair made of a centroidal Voronoi partition P and the corresponding vector of centroids c is locally optimal in the following sense: $\mathcal{H}_{\text{expected}}$ cannot be




Fig. 2. All possible centroidal Voronoi partitions of a uniform 2×5 grid. Assuming all edge weights are w and all vertices have priority 1, then (a) has a cost of $1.2w$, (b) has a cost of $1.1w$, and (c) has a cost of $1.0w$. Only (c) is pairwise-optimal by definition.

reduced by changing either P or c independently. A pairwise-optimal partition achieves this property and adds that for every pair of neighboring robots (i, j) , there does not exist a two-partition of $P_i \cup P_j$ with a lower coverage cost. In other words, positioning the robots at the centroids of a centroidal Voronoi partition (locally) minimizes the expected distance between a task appearing randomly in Q according to relative weights ϕ and the robot who owns the vertex where the task appears. Positioning at the centroids of a pairwise-optimal partition improves performance by reducing the number of sub-optimal solutions which the team might converge to.

III. MODELS, PROBLEM FORMULATION, AND PROPOSED SOLUTION

We aim to partition Q among N robotic agents using only asynchronous, unreliable, short-range communication. In Section III-A we describe the computation, motion, and communication capabilities required of the team of robots, and in Section III-B we formally state the problem we are addressing. In Section III-C we propose our solution, the *Discrete Gossip Coverage Algorithm*, and in III-D we provide an illustration. In Sections III-E and III-F we state the algorithm's convergence and complexity properties.

A. Robot Network Model with Gossip Communication

Our Discrete Gossip Coverage Algorithm requires a team of N robotic agents where each agent $i \in \{1, \dots, N\}$ has the following basic computation and motion capabilities:

- (C1) agent i knows its unique identifier i ;
- (C2) agent i has a processor with the ability to store $G(Q)$ and perform operations on subgraphs of $G(Q)$; and
- (C3) agent i can determine which vertex in Q it occupies and can move at speed v along the edges of $G(Q)$ to any other vertex in Q .

Remark 3.1 (Localization): The localization requirement in (C3) is actually quite loose. Localization is only used for navigation and not for updating partitions, thus limited duration localization errors are not a problem.

The robotic agents are assumed to be able to communicate with each other according to the *range-limited gossip communication model* which is described as follows:

- (C4) given a communication range $r_{\text{comm}} > \max_{e \in E} w_e$, when any two agents reside for some positive duration at a distance $r < r_{\text{comm}}$, they communicate at the sample times of a Poisson process with intensity $\lambda_{\text{comm}} > 0$.

Recall that an homogeneous Poisson process is a widely-used stochastic model for events which occur randomly and

independently in time, where the expected number of events in a period Δ is $\Delta\lambda_{\text{comm}}$.

Remark 3.2 (Communication Model): (1) This communication capability is the minimum necessary for our algorithm, any additional capability can only reduce the time required for convergence. For example, it would be acceptable to have intensity $\lambda(r)$ depend upon the pairwise robot distance in such a way that $\lambda(r) \geq \lambda_{\text{comm}}$ for $r < r_{\text{comm}}$.

(2) We use distances in the graph to model limited range communication. These graph distances are assumed to approximate geodesic distances in the underlying continuous environment and thus path distances for a diffracting wave or moving robot.

B. Problem Statement

Assume that, for all $t \in \mathbb{R}_{\geq 0}$, each agent $i \in \{1, \dots, N\}$ maintains in memory a connected subset $P_i(t)$ of environment Q . Our goal is to design a distributed algorithm that iteratively updates the partition $P(t) = \{P_i(t)\}_{i=1}^N$ while solving the following optimization problem:

$$\min_{P \in \text{Part}_N(Q)} \mathcal{H}_{\text{expected}}(P), \quad (1)$$

subject to the constraints imposed by the robot network model with range-limited gossip communication from Section III-A.

C. The Discrete Gossip Coverage Algorithm

In the design of an algorithm for the minimization problem (1) there are two main questions which must be addressed. First, given the limited communication capabilities in (C4), how should the robots move inside Q to guarantee frequent enough meetings between pairs of robots? Second, when two robots are communicating, what information should they exchange and how should they update their regions?

In this section we introduce the *Discrete Gossip Coverage Algorithm* which, following these two questions, consists of two components:

- (1) the *Random Destination & Wait Motion Protocol*; and
- (2) the *Pairwise Partitioning Rule*.

The concurrent implementation of the Random Destination & Wait Motion Protocol and the Pairwise Partitioning Rule determines the evolution of the positions and dominance subgraphs of the agents as we now formally describe. We start with the Random Destination & Wait Motion Protocol.

Random Destination & Wait Motion Protocol

Each agent $i \in \{1, \dots, N\}$ determines its motion by repeatedly performing the following actions:

- 1: agent i samples a *destination vertex* q_i from a uniform distribution over its dominance subgraph P_i ;
- 2: agent i moves to vertex q_i through the shortest path in P_i connecting the vertex it currently occupies and q_i ; and
- 3: agent i waits at q_i for a duration $\tau > 0$.

If agent i is moving from one vertex to another we say that agent i is in the *moving* state while if agent i is waiting at some vertex we say that it is in the *waiting* state.

Remark 3.3 (Motion Protocol): The motion protocol is designed to ensure frequent enough communication between

pairs of robots. In general, any motion protocol can be used which meets this requirement, so i could select q_i from the boundary of P_i or use some heuristic non-uniform distribution over P_i .

If any two agents i and j reside in two vertices at a graphical distance smaller than r_{comm} for some positive duration, then at the sample times of the corresponding communication Poisson process the two agents exchange sufficient information to update their respective dominance subgraphs P_i and P_j via the Pairwise Partitioning Rule.

Pairwise Partitioning Rule

Assume that at time $t \in \mathbb{R}_{\geq 0}$, agent i and agent j communicate. Without loss of generality assume that $i < j$. Let $P_i(t)$ and $P_j(t)$ denote the current dominance subgraphs of i and j , respectively. Moreover, let t^+ denote the time instant just after t . Then, agents i and j perform the following tasks:

- 1: agent i transmits $P_i(t)$ to agent j and vice-versa
- 2: initialize $W_{a^*} := P_i(t)$, $W_{b^*} := P_j(t)$, $a^* := \text{Cd}(P_i(t))$, $b^* := \text{Cd}(P_j(t))$
- 3: compute $U := P_i(t) \cup P_j(t)$ and an ordered list S of all pairs of vertices in U
- 4: **for** each $(a, b) \in S$ **do**
- 5: compute the sets

$$W_a := \{x \in U : d_U(x, a) \leq d_U(x, b)\}$$

$$W_b := \{x \in U : d_U(x, a) > d_U(x, b)\}$$
- 6: **if** $\mathcal{H}_{\text{one}}(a; W_a) + \mathcal{H}_{\text{one}}(b; W_b) < \mathcal{H}_{\text{one}}(a^*; W_{a^*}) + \mathcal{H}_{\text{one}}(b^*; W_{b^*})$ **then**
- 7: $W_{a^*} := W_a, W_{b^*} := W_b, a^* := a, b^* := b$
- 8: $P_i(t^+) := W_{a^*}, P_j(t^+) := W_{b^*}$

Some remarks are now in order.

Remark 3.4 (Partitioning Rule): (1) The Pairwise Partitioning Rule is designed to find a minimum cost two-partition of U . More formally, if list S and sets W_{a^*} and W_{b^*} for $(a^*, b^*) \in S$ are defined as in the Pairwise Partitioning Rule, then W_{a^*} and W_{b^*} are an optimal two-partition of U .

(2) While the loop in steps 4-7 must run to completion to guarantee that W_{a^*} and W_{b^*} are an optimal two-partition of U , the loop is designed to return an intermediate sub-optimal result if need be. If P_i and P_j change, then $\mathcal{H}_{\text{expected}}$ will decrease and this is enough to ensure eventual convergence.

(3) We make a simplifying assumption in the Pairwise Partitioning Rule that, once two agents communicate, the application of the partitioning rule is instantaneous. We discuss the actual computation time required in Section III-F and some implementation details in Section V.

(4) Notice that simply assigning W_{a^*} to i and W_{b^*} to j can cause the robots to “switch sides” in U . While convergence is guaranteed regardless, switching may be undesirable in some applications. In that case, any smart matching of W_{a^*} and W_{b^*} to i and j may be inserted.

(5) Agents who are not adjacent may communicate but the partitioning rule will not change their regions. Indeed, in this case W_{a^*} and W_{b^*} will not change from $P_i(t)$ and $P_j(t)$.

Some possible modifications and extensions to the algorithm are worth mentioning.

Remark 3.5 (Heterogeneous Robotic Networks): In case the robots have heterogeneous dynamics, line 5 can be




Fig. 3. Simulation of four robots dividing a square environment with obstacles. The boundary of each robots territory is drawn in a different color, the centroid of a territory is drawn with an X, and pairwise communication is drawn with a solid red line. On the left is the initial partition assigned to the robots. The middle frames show two pairwise territory exchanges, with updated territories highlighted with solid colors. The final partition is shown at right.

modified to consider per-robot travel times between vertices. For example, $d_U(x, a)$ could be replaced by the expected time for robot i to travel from a to x while $d_U(x, b)$ would consider robot j .

Remark 3.6 (Coverage and Task Servicing): Here we focus on partitioning territory, but this algorithm can easily be combined with methods to provide a service in Q as in [4]. The agents could split their time between moving to meet their neighbors and update territory, and performing requested tasks in their region.

D. Illustrative Simulation

The simulation in Fig. 3 shows four robots partitioning a square environment with obstacles where the free space is represented by a 12×12 grid. In the initial partition shown in the left panel, the robot in the top right controls most of the environment while the robot in the bottom left controls very little. The robots then move according to the Random Destination & Wait Motion Protocol, and communicate according to range-limited gossip communication model with $r_{comm} = 2.5m$ (four edges in the graph).

The first pairwise territory exchange is shown in the second panel, where the bottom left robot claims some territory from the robot on the top left. A later exchange between the two robots on the top is shown in the next two panels. Notice that the cyan robot in the top right gives away the vertex it currently occupies. In such a scenario, we direct the robot to follow the shortest path in $G(Q)$ to its updated territory before continuing on to a random destination.

After 9 pairwise territory exchanges, the robots reach the pairwise-optimal partition shown at right in Fig. 3. The expected distance between a random vertex and the closest robot decreases from $2.34m$ down to $1.74m$.

E. Convergence Property

The strength of the Discrete Gossip Coverage Algorithm is the possibility of enforcing that a partition will converge to a pairwise-optimal partition through pairwise territory exchange. In Theorem 3.7 we summarize this convergence property, with proofs given in Section IV.

Theorem 3.7 (Convergence Property): Consider a network of N robotic agents endowed with computation and motion capacities (C1), (C2), (C3), and communication capacities (C4). Assume the agents implement the *Discrete Gossip Coverage*

Algorithm consisting of the concurrent implementation of the *Random Destination & Wait Motion Protocol* and the *Pairwise Partitioning Rule*. Then,

- (i) the partition $P(t)$ remains connected and is described by $P : \mathbb{R}_{\geq 0} \rightarrow \text{Part}_N(Q)$, and
- (ii) $P(t)$ converges almost surely in finite time to a pairwise-optimal partition.

Remark 3.8 (Optimality of Solutions): By definition, a pairwise-optimal partition is optimal in that $\mathcal{H}_{\text{expected}}$ can not be improved by changing only two regions in the partition.

Remark 3.9 (Generalizations): For simplicity we assume uniform robot speeds, communication processes, and waiting times. An extension to non-uniform processes would be straightforward.

F. Complexity Properties and Discussion

In this subsection we explore the computational requirements of the Discrete Gossip Coverage Algorithm, and make some comments on implementation. Cost function $\mathcal{H}_{\text{one}}(h; P_i)$ is the sum of the distances between h and all other vertices in P_i . This computation of one-to-all distances is the core computation of the algorithm. For most graphs of interest the total number of edges $|E|$ is proportional to $|Q|$, so we will state bounds on this computation in terms of $|P_i|$. Computing one-to-all distances requires one of the following:

- if all edge weights in $G(Q)$ are the same (e.g., for a graph from an occupancy grid), a breadth-first search approach can be used which requires $\mathcal{O}(|P_i|)$ in time and memory;
- otherwise, Dijkstra's algorithm must be used which requires $\mathcal{O}(|P_i| \log(|P_i|))$ in time and $\mathcal{O}(|P_i|)$ in memory.

Let $\mathbb{D}(P_i)$ be the time to compute one-to-all distances in P_i , then computing $\mathcal{H}_{\text{one}}(h; P_i)$ requires $\mathcal{O}(\mathbb{D}(P_i))$ in time.

Proposition 3.10 (Complexity Properties): The motion protocol requires $\mathcal{O}(|P_i|)$ in memory, and $\mathcal{O}(\mathbb{D}(P_i))$ in computation time. The partitioning rule requires $\mathcal{O}(|P_i| + |P_j|)$ in communication bandwidth between robots i and j , $\mathcal{O}(|P_i| + |P_j|)$ in memory, and can run in any time.

Proof: We first prove the claims for the motion protocol. Step 2 is the only non-trivial step and requires finding a shortest path in P_i , which is equivalent to computing one-to-all distances from the robot's current vertex. Hence, it requires $\mathcal{O}(\mathbb{D}(P_i))$ in time and $\mathcal{O}(P_i)$ in memory.

We now prove the claims for the partitioning rule. In step 1, robots i and j transmit their subgraphs to each other,

which requires $\mathcal{O}(|P_i| + |P_j|)$ in communication bandwidth. For step 3, the robots determine $U := P_i \cup P_j$, which requires $\mathcal{O}(|P_i| + |P_j|)$ in memory to store. Step 4 is the start of a loop which executes $\mathcal{O}(|U|^2)$ times, affecting the time complexity of steps 5, 6 and 7. Step 5 requires two computations of one-to-all distances in U which each take $\mathcal{O}(\mathbb{D}(U))$. Step 6 involves four computations of \mathcal{H}_{one} over different subsets of U , however those for W_{a^*} and W_{b^*} can be stored from previous computation. Since W_a and W_b are strict subsets of U , step 5 takes longer than step 6. Step 7 is trivial, as is step 8. The total time complexity of the loop is thus $\mathcal{O}(|U|^2 \mathbb{D}(U))$.

However, the loop in steps 4-7 can be truncated after any number of iterations. While it must run to completion to guarantee that W_{a^*} and W_{b^*} are an optimal two-partition of U , the loop is designed to return an intermediate sub-optimal result if need be. If P_i and P_j change, then $\mathcal{H}_{\text{expected}}$ will decrease. Our convergence result will hold provided that all elements of S are eventually checked if P_i and P_j do not change. Thus, the partitioning rule can run in any time with each iteration requiring $\mathcal{O}(\mathbb{D}(U))$. ■

All of the computation and communication requirements in Proposition 3.10 are independent of the number of robots and scale with the size of a robot's partition, meaning the Discrete Gossip Coverage Algorithm can easily scale up for large teams of robots in large environments.

IV. CONVERGENCE PROOFS

This section is devoted to proving the two statements in Theorem 3.7. The proof that the Pairwise Partitioning Rule maps a connected N -partition into a connected N -partition is straightforward. The proof of convergence is more involved and is based on the application of Lemma A.1 in Appendix A to the Discrete Gossip Coverage Algorithm. Lemma A.1 establishes strong convergence properties for a particular class of set valued maps (set-valued maps are briefly reviewed in Appendix A).

We start by proving that the Pairwise Partitioning Rule is well-posed in the sense that it maintains a connected partition.

Proof of Theorem 3.7 statement (i): To prove the statement we need to show that $P(t^+)$ satisfies points (i) through (iv) of Definition 2.2. From the definition of the Pairwise Partitioning Rule, we have that $P_i(t^+) \cup P_j(t^+) = P_i(t) \cup P_j(t)$ and $P_i(t^+) \cap P_j(t^+) = \emptyset$. Moreover, since $a^* \in P_i(t^+)$ and $b^* \in P_j(t^+)$, it follows that $P_i(t^+) \neq \emptyset$ and $P_j(t^+) \neq \emptyset$. These observations imply the validity of points (i), (ii), and (iii) for $P(t^+)$. Finally, we must show that $P_i(t^+)$ and $P_j(t^+)$ are connected, i.e., $P(t^+)$ also satisfies point (iv). To do so we show that, given $x \in W_{a^*}$, any shortest path in $P_i(t) \cup P_j(t)$ connecting x to a^* completely belongs to W_{a^*} . We proceed by contradiction. Let s_{x,a^*} denote a shortest path in $P_i(t) \cup P_j(t)$ connecting x to a^* and let us assume that there exists $m \in s_{x,a^*}$ such that $m \in W_{b^*}$. For m to be in W_{b^*} means that $d_{P_i(t) \cup P_j(t)}(m, b^*) < d_{P_i(t) \cup P_j(t)}(m, a^*)$. This implies that

$$\begin{aligned} d_{P_i \cup P_j}(x, b^*) &\leq d_{P_i \cup P_j}(m, b^*) + d_{P_i \cup P_j}(x, m) \\ &< d_{P_i \cup P_j}(m, a^*) + d_{P_i \cup P_j}(x, m) \\ &= d_{P_i \cup P_j}(x, a^*). \end{aligned}$$

This is a contradiction for $x \in W_{a^*}$. Similar considerations hold for W_{b^*} . ■

The rest of this section is dedicated to proving convergence. Our first step is to show that the evolution determined by the Discrete Gossip Coverage Algorithm can be seen as a set-valued map. To this end, for any pair of robots $(i, j) \in \{1, \dots, N\}^2$, $i \neq j$, we define the map $T_{ij} : \text{Part}_N(Q) \rightarrow \text{Part}_N(Q)$ by

$$T_{ij}(P) = (P_1, \dots, \widehat{P}_i, \dots, \widehat{P}_j, \dots, P_N),$$

where $\widehat{P}_i = W_{a^*}$ and $\widehat{P}_j = W_{b^*}$.

If at time $t \in \mathbb{R}_{\geq 0}$ the pair (i, j) and no other pair of robots perform an iteration of the Pairwise Partitioning Rule, then the dynamical system on the space of partitions is described by

$$P(t^+) = T_{ij}(P(t)). \quad (2)$$

We define the set-valued map $T : \text{Part}_N(Q) \rightrightarrows \text{Part}_N(Q)$ as

$$T(P) = \{T_{ij}(P) \mid (i, j) \in \{1, \dots, N\}^2, i \neq j\}. \quad (3)$$

Observe that (2) can then be rewritten as $P(t^+) \in T(P(t))$.

The next two Propositions state facts whose validity is ensured by Lemma B.1 of Appendix B which states a key property of the Random Destination & Wait Motion Protocol.

Proposition 4.1 (Persistence of Exchanges): Consider N robots implementing the Discrete Gossip Coverage Algorithm. Then, there almost surely exists an increasing sequence of time instants $\{t_k\}_{k \in \mathbb{Z}_{\geq 0}}$ such that $P(t_k^+) = T_{ij}(P(t_k))$ for some $(i, j) \in \mathcal{E}(P(t_k))$.

Proof: The proof follows directly from Lemma B.1 which implies that the time between two consecutive pairwise communications is almost surely finite. ■

The existence of time sequence $\{t_k\}_{k \in \mathbb{Z}_{\geq 0}}$ allows us to express the evolution generated by the Discrete Gossip Coverage Algorithm as a discrete time process. Let $P(k) := P(t_k)$ and $P(k+1) := P(t_k^+)$, then

$$P(k+1) \in T(P(k))$$

where $T : \text{Part}_N(Q) \rightrightarrows \text{Part}_N(Q)$ is defined as in (3).

Given $k \in \mathbb{Z}_{\geq 0}$, let \mathcal{I}_k denote the information which completely characterizes the state of Discrete Gossip Coverage Algorithm just after the k -th iteration of the partitioning rule, i.e., at time t_{k-1}^+ . Specifically, \mathcal{I}_k contains the information related to the partition $P(k)$, the positions of the robots at t_{k-1}^+ , and whether each robot is in the *waiting* or *moving* state at t_{k-1}^+ . The following result characterizes the probability that, given \mathcal{I}_k , the $(k+1)$ -th iteration of the partitioning rule is governed by any of the maps T_{ij} , $(i, j) \in \mathcal{E}(P(k))$.

Proposition 4.2 (Probability of Communication): Consider a team of N robots with capacities (C1), (C2), (C3), and (C4) implementing the Discrete Gossip Coverage Algorithm. Then, there exists a real number $\bar{\pi} \in (0, 1)$, such that, for any $k \in \mathbb{Z}_{\geq 0}$ and $(i, j) \in \mathcal{E}(P(k))$

$$\mathbb{P}[P(k+1) = T_{ij}(P(k)) \mid \mathcal{I}_k] \geq \bar{\pi}.$$

Proof: Assume that at time \bar{t} one pair of robots communicates. Given a pair $(\bar{i}, \bar{j}) \in \mathcal{E}(P(\bar{t}))$, we must find a lower bound for the probability that (\bar{i}, \bar{j}) is the communicating pair.

Since all the Poisson communication processes have the same intensity, the distribution of the chance of communication is uniform over the pairs which are “able to communicate,” i.e., closer than r_{comm} to each other. Thus, we must only show that (\bar{i}, \bar{j}) has a positive probability of being able to communicate at time \bar{t} , which is equivalent to showing that (\bar{i}, \bar{j}) is able to communicate for a positive fraction of time with positive probability. The proof of Lemma B.1 implies that with probability at least $\alpha/(1 - e^{-\lambda_{\text{comm}}\tau})$ any pair in $\mathcal{E}(P(\bar{t}))$ is able to communicate for a fraction of time not smaller than $\frac{\tau}{\Delta}$, where α and Δ are defined in the proof of Lemma B.1. Hence the result follows. ■

The property in Proposition 4.2 can also be formulated as follows. Let $\sigma : \mathbb{Z}_{\geq 0} \rightarrow \{(i, j) \in \{1, \dots, N\}^2, i \neq j\}$ be the stochastic process such that $\sigma(k)$ is the communicating pair at time k . Then, the sequence of pairs of robots performing the partitioning rule at time instants $\{t_k\}_{k \in \mathbb{Z}_{\geq 0}}$ can be seen as a realization of the process σ , which satisfies

$$\mathbb{P}[\sigma(k+1) = (i, j) \mid \sigma(k)] \geq \bar{\pi} \quad (4)$$

for all $(i, j) \in \mathcal{E}(P(k))$.

Next we show that the cost function decreases whenever the application of T from (3) changes the territory partition. This fact is a key ingredient to apply Lemma A.1.

Lemma 4.3 (Decreasing Cost Function): Let $P \in \text{Part}_N(Q)$ and let $P^+ \in T(P)$. If $P^+ \neq P$, then $\mathcal{H}_{\text{expected}}(P^+) < \mathcal{H}_{\text{expected}}(P)$.

Proof: Without loss of generality assume that (i, j) is the pair executing the Pairwise Partitioning Rule. Then

$$\begin{aligned} \mathcal{H}_{\text{expected}}(P^+) - \mathcal{H}_{\text{expected}}(P) &= \mathcal{H}_{\text{one}}(\text{Cd}(P_i^+); P_i^+) + \mathcal{H}_{\text{one}}(\text{Cd}(P_j^+); P_j^+) \\ &\quad - \mathcal{H}_{\text{one}}(\text{Cd}(P_i); P_i) - \mathcal{H}_{\text{one}}(\text{Cd}(P_j); P_j). \end{aligned}$$

According to the definition of the Pairwise Partitioning Rule we have that if $P_i^+ \neq P_i$, $P_j^+ \neq P_j$, then

$$\begin{aligned} \mathcal{H}_{\text{one}}(\text{Cd}(P_i^+); P_i^+) + \mathcal{H}_{\text{one}}(\text{Cd}(P_j^+); P_j^+) & \\ \leq \mathcal{H}_{\text{one}}(a^*; P_i^+) + \mathcal{H}_{\text{one}}(b^*; P_j^+) & \\ < \mathcal{H}_{\text{one}}(\text{Cd}(P_i); P_i) + \mathcal{H}_{\text{one}}(\text{Cd}(P_j); P_j) & \end{aligned}$$

from which the statement follows. ■

We now complete the proof of the main result, Theorem 3.7.

Proof of Theorem 3.7 statement (ii): Note that the algorithm evolves in a finite space of partitions, and by Theorem 3.7 statement (i), the set $\text{Part}_N(Q)$ is strongly positively invariant. This fact implies that assumption (i) of Lemma A.1 is satisfied. From Lemma 4.3 it follows that assumption (ii) is also satisfied, with $\mathcal{H}_{\text{expected}}$ playing the role of the function U . Finally, the property in (4) is equivalent to the property of *persistent random switches* stated in Assumption (iii) of Lemma A.1, for the special case $h = 1$. Hence, we are in the position to apply Lemma A.1 and conclude convergence in finite-time to an element of the intersection of the equilibria of the maps T_{ij} , which by definition is the set of the pairwise-optimal partitions. ■

V. EXPERIMENTAL METHODS & RESULTS

To demonstrate the utility and study practical issues of the Discrete Gossip Coverage Algorithm, we implemented it using the open-source Player/Stage robot control system [23] and the Boost Graph Library (BGL) [24]. All results presented here were generated using Player 2.1.1, Stage 2.1.1, and BGL 1.34.1. To compute distances in uniform edge weight graphs we extended the BGL breadth-first search routine with a distance recorder event visitor.

A. Large-scale Simulation

To evaluate the performance of our gossip coverage algorithm with larger teams, we tested 30 simulated robots partitioning a map representing a $350m \times 225m$ portion of campus at the University of California at Santa Barbara. As shown in Fig. 4, the robots are tasked with providing coverage of the open space around some of the buildings on campus, a space which includes a couple open quads, some narrower passages between buildings, and a few dead-end spurs. For this large environment the simulated robots are $2m$ on a side and can move at $3.0 \frac{m}{s}$. Each territory cell is $3m \times 3m$.

In this simulation we handle communication and partitioning as follows. The communication range is set to $30m$ (10 edges in the graph) with $\lambda_{\text{comm}} = 0.3 \frac{\text{comm}}{s}$. The robots wait at their destination vertices for $\tau = 3.5s$. This value for τ was chosen so that on average one quarter of the robots are waiting at any moment. Lower values of τ mean the robots are moving more of the time and as a result more frequently miss connections, while for higher τ the robots spend more time stationary which also reduces the rate of convergence. With the goal of improving communication, we implemented a minor modification to the motion protocol: each robot picks its random destination from the cells forming the open boundary¹ of its territory. In our implementation, the full partitioning loop may take 5 seconds for the largest initial territories in Fig. 4. We chose to stop the loop after a quarter second for this simulation to verify the anytime computation claim.

The 30 robots start clustered in the center of the map between Engineering II and Broida Hall, and an initial Voronoi partition is generated from these starting positions. This initial partition is shown on the left in Fig. 4 with the robots positioned at the centroids of their starting regions. The initial partition has a cost of $37.1m$. The team spends about 27 minutes moving and communicating according to the Discrete Gossip Coverage Algorithm before settling on the final partition on the right of Fig. 4. The coverage cost of the final equilibrium improved by 54% to $17.1m$. Visually, the final partition is also dramatically more uniform than the initial condition. This result demonstrates that the algorithm is effective for large teams in large non-convex environments.

Fig. 5 shows the evolution of $\mathcal{H}_{\text{expected}}$ during the simulation. The largest cost improvements happen early when the robots that own the large territories on the left and right of the map communicate with neighbors with much smaller territories.

¹The open boundary of P_i is the set of vertices in P_i which are adjacent to at least one vertex owned by another agent.

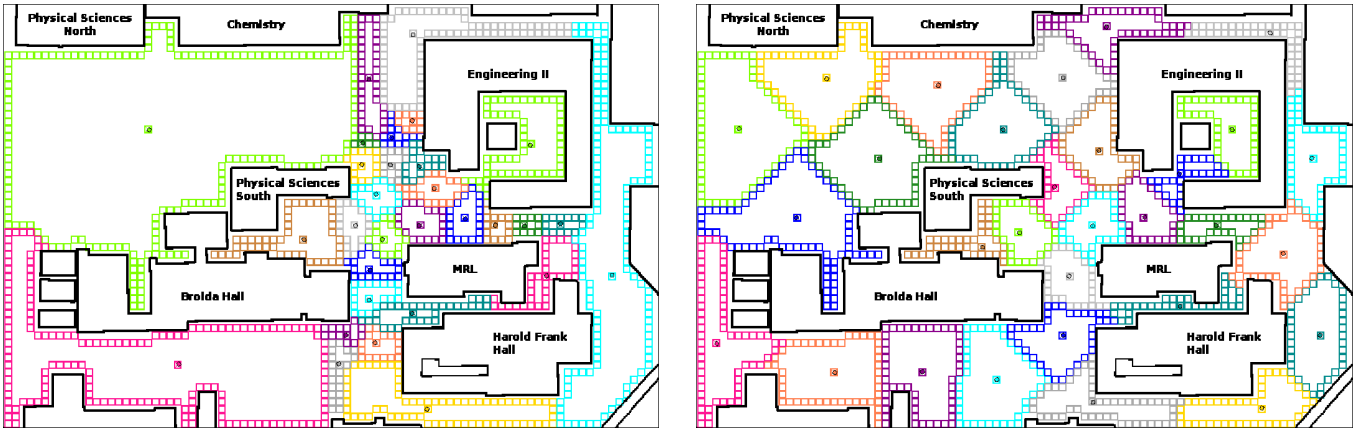


Fig. 4. Images of starting and final partitions for a simulation with 30 robots providing coverage of a portion of campus at UCSB.




Fig. 5. Graph of the cost $\mathcal{H}_{\text{expected}}$ over time for the simulation in Fig. 4.

These big territory changes then propagate through the network as the robots meet and are pushed and pulled towards a lower cost partition.

B. Implementation Details

We conducted an experiment to test the algorithm using three physical robots in our lab, augmented by six simulated robots in a synthetic environment extending beyond the lab. Our lab space is $11.3m$ on a side and is represented by the upper left portion of the territory maps in Fig. 7. The territory graph loops around a center island of desks. We extended the lab space through three connections into a simulated environment around the lab, producing a $15.9m \times 15.9m$ environment. The map of the environment was specified with a $0.15m$ bitmap which we overlaid with a $0.6m$ resolution occupancy grid representing the free territory for the robots to cover. The result is a lattice-like graph with all edge weights equal to $0.6m$. The $0.6m$ resolution was chosen so that our physical robots would fit easily inside a cell.

Additional details of our implementation are as follows.

Robot hardware: We use Erratic mobile robots from Videre Design, as shown in Fig. 6. The vehicle platform has a roughly square footprint ($40cm \times 37cm$), with two differential drive wheels and a single rear caster. Each robot carries an onboard computer with a 1.8Ghz Core 2 Duo processor, 1 GB of memory, and 802.11g wireless communication. For navigation and localization, each robot is equipped with a Hokuyo URG-04LX laser rangefinder. The rangefinder scans 683 points over 240° at $10Hz$ with a range of 5.6 meters.



Fig. 6. Erratic mobile robot with URG-04LX laser rangefinder.

Experiment setup: Our mixed physical and virtual robot experiments are run from a central computer which is attached to a wireless router so it can communicate with the physical robots. The central computer creates a simulated world using Stage which mirrors and extends the real space in which the physical robots operate. The central computer also simulates the virtual members of the robot team. These virtual robots are modeled off of our hardware: they are differential drive with the same geometry as the Erratic platform and use simulated Hokuyo URG-04LX rangefinders.

Localization: We use the `amcl` driver in Player which implements Adaptive Monte-Carlo Localization [25]. The physical robots are provided with a map of our lab with a $15cm$ resolution and told their starting pose within the map. We set an initial pose standard deviation of $0.9m$ in position and 12° in orientation, and request localization updates using 50 of the sensor's range measurements for each change of $2cm$ in position or 2° in orientation reported by the robot's odometry system. We then use the most likely pose estimate output by `amcl` as the location of the robot. For simplicity and reduced computational demand, we allow the virtual robots access to perfect localization information.

Motion Protocol: Each robot continuously executes the Random Destination & Wait Motion Protocol, with navigation handled by the `snd` driver in Player which implements Smooth Nearness Diagram navigation [26]. For `snd` we set the robot radius parameter to $22cm$, obstacle avoidance distance to $0.7m$, and maximum speeds to $0.4 \frac{m}{s}$ and $40 \frac{c}{s}$. The `snd` driver is a local obstacle avoidance planner, so we feed it a series of waypoints every couple meters along paths found in $G(Q)$.




Fig. 7. Each column contains a territory map and the corresponding overhead camera image for a step of the hardware-in-the-loop simulation. The position of the camera in the environment is shown with a camera icon in the territory map. The physical robots are numbered 1, 2, and 3 and have the orange, blue, and lime green partitions. Their positions in each territory map are indicated with numbered circles.

We consider a robot to have achieved its target location when it is within 20cm and it will then wait for $\tau = 3.5\text{s}$. For the physical robots the motion protocol and navigation processes run on board, while there are separate threads for each virtual robot on the central computer.

Communication and Partitioning: As the robots move, a central process monitors their positions and simulates the range-limited gossip communication model between both real and virtual robots. We set $r_{\text{comm}} = 2.5\text{m}$ and $\lambda_{\text{comm}} = 0.3 \frac{\text{comm}}{\text{s}}$. These parameters were chosen so that the robots would be likely to communicate when separated by at most four edges, but would also sometimes not connect despite being close. When this process determines two robots should communicate, it informs the robots who then perform the Pairwise Partitioning Rule. Our pairwise communication implementation is blocking: if robot i is exchanging territory with j , then it informs the match making process that it is unavailable until the exchange is complete.

C. Hardware-in-the-Loop Simulation

The results of our experiment with three physical robots and six simulated robots are shown in Figs. 7 and 8. The left column in Fig. 7 shows the starting positions of the team of robots, with the physical robots, labeled 1, 2, and 3, lined up in a corner of the lab and the simulated robots arrayed around them. The starting positions are used to generate the initial Voronoi partition of the environment. The physical robots own the orange, blue, and lime green territories in the upper left quadrant. We chose this initial configuration to have a high coverage cost, while ensuring that the physical robots will remain in the lab as the partition evolves.




Fig. 8. Evolution of cost functions during the experiment in Fig. 7. The total cost $\mathcal{H}_{\text{expected}}$ is shown above in black, while \mathcal{H}_{one} for each robot is shown below in the robot's color.

In the middle column, robots 1 and 2 have met along their shared border and are exchanging territory. In the territory map, the solid red line indicates 1 and 2 are communicating and their updated territories are drawn with solid orange and blue, respectively. The camera view confirms that the two robots have met on the near side of the center island of desks.

The final partition at right in Fig. 7 is reached after $9\frac{1}{2}$ minutes. All of the robots are positioned at the centroids of their final territories. The three physical robots have gone from a cluster in one corner of the lab to a more even spread around the space. Fig. 8 shows the evolution of the cost function

$\mathcal{H}_{\text{expected}}$ as the experiment progresses, including the costs for each robot. As expected, the total cost never increases and the disparity of costs for the individual robots shrinks over time until settling at a pairwise-optimal partition.

In this experiment the hardware challenges of sensor noise, navigation, and uncertainty in position were efficiently handled by the `amcl` and `snd` drivers. The coverage algorithm assumed the role of a higher-level planner, taking in position data from `amcl` and directing `snd`. By far the most computationally demanding component was `amcl`, but the position hypotheses from `amcl` are actually unnecessary: our coverage algorithm only requires knowledge of the vertex a robot occupies. If a less intensive localization method is available, the algorithm could run on robots with significantly lower compute power.

D. Comparative analysis

In this subsection we present a numerical comparison of the performance of the Discrete Gossip Coverage Algorithm and the following two Lloyd-type algorithms.

Decentralized Lloyd Algorithm: This method is from [11] and [9], we describe it here for convenience. At each discrete time instant $t \in \mathbb{Z}_{\geq 0}$, each robot i performs the following tasks: (1) i transmits its position and receives the positions of all adjacent robots; (2) i computes its Voronoi region P_i based on the information received; and (3) i moves to $\text{Cd}(P_i)$.

Gossip Lloyd Algorithm: This method is from [19]. It is a gossip algorithm, and so we have used the same communication model and the Random Destination & Wait Motion Protocol to create meetings between robots. Say robots i and j meet at time t , then the pairwise Lloyd partitioning rule works as follows: (1) robot i transmits $P_i(t)$ to j and vice versa; (2) both robots determine $U = P_i(t) \cup P_j(t)$; (3) robot i sets $P_i(t^+)$ to be its Voronoi region of U based on $\text{Cd}(P_i(t))$ and $\text{Cd}(P_j(t))$, and j does the equivalent.

For both Lloyd algorithms we use the same tie breaking rule when creating Voronoi regions as is present in the Pairwise Partitioning Rule: ties go to the robot with the lowest index.

Our first numerical result uses a Monte Carlo probability estimation method from [27] to place probabilistic bounds on the performance of the two gossip algorithms. Recall that the Chernoff bound describes the minimum number of random samples K required to reach a certain level of accuracy in a probability estimate from independent Bernoulli tests. For an accuracy $\epsilon \in (0, 1)$ and confidence $1 - \eta \in (0, 1)$, the number of samples is given by $K \geq \frac{1}{2\epsilon^2} \log \frac{2}{\eta}$. For $\eta = 0.01$ and $\epsilon = 0.1$, at least 116 samples are required.

Figure 9 shows both the initial territory partition of the extended laboratory environment used and also a histogram of the final results for the following Monte Carlo test. The environment and robot motion models used are described in Section V-B. Starting from the indicated initial condition, we ran 116 simulations of both gossip algorithms. The randomness in the test comes from the sequence of pairwise communications. These sequences were generated using: (1) the Random Destination & Wait Motion Protocol with q_i sampled uniformly from the open boundary of P_i and $\tau = 3.5s$;




Fig. 9. Initial partition and histogram of final costs for a Monte Carlo test comparing the Discrete Gossip Coverage Algorithm (black bars), Gossip Lloyd Algorithm (gray bars), and Decentralized Lloyd Algorithm (red dashed line). For the gossip algorithms, 116 simulations were performed with different sequences of pairwise communications. The Decentralized Lloyd Algorithm is deterministic given an initial condition so only one final cost is shown.

and (2) the range-limited gossip communication model with $r_{\text{comm}} = 2.5m$ and $\lambda_{\text{comm}} = 0.3 \frac{\text{comm}}{s}$.

The cost of the initial partition in Fig. 9 is $5.48m$, while the best known partition for this environment has a cost of just under $2.18m$. The histogram in Fig. 9 shows the final equilibrium costs for 116 simulations of the Discrete Gossip Coverage Algorithm (black) and the Gossip Lloyd Algorithm (gray). It also shows the final cost using the Decentralized Lloyd Algorithm (red dashed line), which is deterministic from a given initial condition. The histogram bins have a width of $0.10m$ and start from $2.17m$. For the Discrete Gossip Coverage Algorithm, 105 out of 116 trials reach the bin containing the best known partition and the mean final cost is $2.23m$. The Gossip Lloyd Algorithm reaches the lowest bin in only 5 of 116 trials and has a mean final cost of $2.51m$. The Decentralized Lloyd Algorithm settles at $2.48m$. Our new gossip algorithm requires an average of 96 pairwise communications to reach an equilibrium, whereas gossip Lloyd requires 126.

Based on these results, we can conclude with 99% confidence that there is at least an 80% probability that 9 robots executing the Discrete Gossip Coverage Algorithm starting from the initial partition shown in Fig. 9 will reach a pairwise-optimal partition which has a cost within 4% of the best known cost. We can further conclude with 99% confidence that the




Fig. 10. Histograms of final costs from 10 Monte Carlo tests using random initial conditions in the environment shown in Fig. 9 comparing Discrete Gossip Coverage Algorithm (black bars), Gossip Lloyd Algorithm (gray bars), and Decentralized Lloyd Algorithm (red dashed line). For the gossip algorithms, 116 simulations were performed with different sequences of pairwise communications. The Decentralized Lloyd Algorithm is deterministic given an initial condition so only one final cost is shown. The initial cost for each test is drawn with the green dashed line.

Gossip Lloyd Algorithm will settle more than 4% above the best known cost at least 86% of the time starting from this initial condition.

Figure 10 compares final cost histograms for 10 different initial conditions for the same environment and parameters as described above. Each initial condition was created by selecting unique starting locations for the robots uniformly at random and using these locations to generate an initial Voronoi partition. The initial cost for each test is shown with the green dashed line. In 9 out of 10 tests the Discrete Gossip Coverage Algorithm reaches the histogram bin with the best known partition in at least 112 of 116 trials. The two Lloyd methods get stuck in sub-optimal centroidal Voronoi partitions more than 4% away from the best known partition in more than

half the trials in 7 of 10 tests.

VI. CONCLUSION

We have presented a novel distributed partitioning and coverage control algorithm which requires only unreliable short-range communication between pairs of robots and works in non-convex environments. The classic Lloyd approach to coverage optimization involves iteration of separate centering and Voronoi partitioning steps. For gossip algorithms, however, this separation is unnecessary computationally and we have shown that improved performance can be achieved without it. Our new Discrete Gossip Coverage Algorithm provably converges to a subset of the set of centroidal Voronoi partitions which we labeled pairwise-optimal partitions. Through numerical comparisons we demonstrated that this new subset of solutions avoids many of the local minima in which Lloyd-type algorithms can get stuck.

Our vision is that this partitioning and coverage algorithm will form the foundation of a distributed task servicing setup for teams of mobile robots. The robots would split their time between servicing tasks in their territory and moving to contact their neighbors and improve the coverage of the space. Our convergence results only require sporadic improvements to the cost function, affording flexibility in robot behaviors and capacities, and offering the ability to handle heterogeneous robotic networks. In the bigger picture, this paper demonstrates the potential of gossip communication in distributed coordination algorithms. There appear to be many other problems where this realistic and minimal communication model could be fruitfully applied.

APPENDIX A

For completeness we present a convergence result for set-valued algorithms on finite state spaces, which can be recovered as a direct consequence of [18, Theorem 4.5].

Given a set X , a set-valued map $T : X \rightrightarrows X$ is a map which associates to an element $x \in X$ a subset $Z \subset X$. A set-valued map is non-empty if $T(x) \neq \emptyset$ for all $x \in X$. Given a non-empty set-valued map T , an evolution of the dynamical system associated to T is a sequence $\{x_n\}_{n \in \mathbb{Z}_{\geq 0}} \subset X$ where $x_{n+1} \in T(x_n)$ for all $n \in \mathbb{Z}_{\geq 0}$. A set $W \subset X$ is *strongly positively invariant* for T if $T(w) \subset W$ for all $w \in W$.

Lemma A.1 (Persistent random switches imply convergence): Let (X, d) be a finite metric space. Given a collection of maps $T_1, \dots, T_m : X \rightarrow X$, define the set-valued map $T : X \rightrightarrows X$ by $T(x) = \{T_1(x), \dots, T_m(x)\}$. Given a stochastic process $\sigma : \mathbb{Z}_{\geq 0} \rightarrow \{1, \dots, m\}$, consider an evolution $\{x_n\}_{n \in \mathbb{Z}_{\geq 0}}$ of T satisfying $x_{n+1} = T_{\sigma(n)}(x_n)$. Assume that:

- (i) there exists a set $W \subseteq X$ that is strongly positively invariant for T ;
- (ii) there exists a function $U : W \rightarrow \mathbb{R}$ such that $U(w') < U(w)$, for all $w \in W$ and $w' \in T(w) \setminus \{w\}$; and
- (iii) there exist $p \in (0, 1)$ and $k \in \mathbb{N}$ such that, for all $i \in \{1, \dots, m\}$ and $n \in \mathbb{Z}_{\geq 0}$, there exists $h \in \{1, \dots, k\}$ such that $\mathbb{P}[\sigma(n+h) = i \mid \sigma(n), \dots, \sigma(1)] \geq p$.

For $i \in \{1, \dots, m\}$, let F_i be the set of fixed points of T_i in W , i.e., $F_i = \{w \in W \mid T_i(w) = w\}$. If $x_0 \in W$, then the

evolution $\{x_n\}_{n \in \mathbb{Z}_{\geq 0}}$ converges almost surely in finite time to an element of the set $(F_1 \cap \dots \cap F_m)$, i.e., there exists almost surely $\tau \in \mathbb{N}$ such that, for some $\bar{x} \in (F_1 \cap \dots \cap F_m)$, $x_n = \bar{x}$ for $n \geq \tau$.

APPENDIX B

This Appendix proves a property of the Random Destination & Wait Motion Protocol which is needed to show the persistence of pairwise exchanges.

Lemma B.1: Consider N robots implementing the Discrete Gossip Coverage Algorithm starting from an arbitrary $P \in \text{Part}_N(Q)$. Consider $t \in \mathbb{R}_{\geq 0}$ and let $P(t)$ denote the partition at time t . Assume that at time t no two robots are communicating. Then, there exist $\Delta > 0$ and $\alpha \in (0, 1)$, independent of $P(t)$ and the positions and states of the robots at time t , such that, for every $(i, j) \in \mathcal{E}(P(t))$, $\mathbb{P}[(i, j) \text{ communicate within } (t, t + \Delta)] \geq \alpha$.

Proof: To begin, we define two useful quantities. Let $\mathcal{S}(Q) := \max_{P \in \text{Part}_N(Q)} \max_{P_i \in P} \max_{h, k \in P_i} d_{P_i}(h, k)$ be a pseudo-diameter for Q , and then choose $\Delta := 2 \frac{\mathcal{S}(Q)}{v} + 2\tau$. We fix a pair $(i, j) \in \mathcal{E}(P)$, and pick adjacent vertices $a \in P_i$, $b \in P_j$.

Our goal is to lower bound the probability that i and j will communicate within the interval $(t, t + \Delta)$. To do so we construct *one* sequence of events of positive probability which enables such communication. Consider the following situation: i is in the *moving* state and needs time t_i to reach its destination q_i , whereas robot j is in the *waiting* state at vertex q_j and must wait there for time $\tau_j \leq \tau$. We denote by $t(a)$ (resp. $t(b)$) the time needed for i (resp. j) to travel from q_i (resp. q_j) to a (resp. b). Let E_i be the event such that i performs the following actions in $(t, t + \Delta)$ without communicating with any robot $k \neq j$:

- (i) i reaches q_i and waits at q_i for the duration τ ; and
- (ii) i chooses vertex a as its next destination and then stays at a for at least $\Delta - t(a) - t_i - \tau$.

Let E_j be the event such that j performs the following actions in $(t, t + \Delta)$ without communicating with any $k \neq i$:

- (i) j waits at q_j for the duration τ_j ; and
- (ii) j chooses vertex b as its next destination and then stays at b for at least $\Delta - t(b) - \tau_j$.

Let $E_{ij} = E_i \cap E_j$.

Next, we lower bound the probability that event E_i occurs. Recall the definition of λ_{comm} from Sec. III-A. Since a robot can have at most $N - 1$ neighbors, the probability that (i) of E_i happens is lower bounded by $e^{-\lambda_{\text{comm}}\tau N}$. For (ii), the probability that i chooses a is $1/|P_i|$, which is lower bounded by $1/|Q|$. Then, in order to spend at least $(\Delta - t(a) - t_i - \tau)$ at a , i must choose a for $\lceil \frac{\Delta - t(a) - t_i - \tau}{\tau} \rceil$ consecutive times. Finally, the probability that during this interval i will not communicate with any robot other than j is lower bounded by $e^{-\lambda_{\text{comm}}\Delta(N-2)}$. The probability that (ii) occurs is thus lower bounded by $(1/|Q|)^{\lceil \frac{\Delta}{\tau} \rceil} e^{-\lambda_{\text{comm}}\Delta N}$. Combining the bounds for (i) and (ii), it follows that

$$\mathbb{P}[E_i] \geq \left(\frac{1}{|Q|}\right)^{\lceil \frac{\Delta}{\tau} \rceil} e^{-\lambda_{\text{comm}}(\Delta + \tau)N}.$$

The same lower bound holds for $\mathbb{P}[E_j]$, meaning that

$$\mathbb{P}[E_{ij}] = \mathbb{P}[E_i] \mathbb{P}[E_j] \geq \left(\frac{1}{|Q|}\right)^{2\lceil \frac{\Delta}{\tau} \rceil} e^{-2\lambda_{\text{comm}}(\Delta + \tau)N}.$$

If event E_{ij} occurs, then robots i and j will be at adjacent vertices for an amount of time during the interval $(t, t + \Delta)$ equal to $\min\{\Delta - t(a) - t_i - \tau, \Delta - t(b) - \tau_j\}$. Since $t(a)$ and $t(b)$ are no more than $\frac{\mathcal{S}(Q)}{v}$, we can conclude that i and j will be within r_{comm} for at least τ . Conditioned on E_{ij} occurring, the probability that i and j communicate in $(t, t + \Delta)$ is lower bounded by $1 - e^{-\lambda_{\text{comm}}\tau}$. A suitable choice for α from the statement of the Lemma is thus

$$\alpha = \left(\frac{1}{|Q|}\right)^{2\lceil \frac{\Delta}{\tau} \rceil} e^{-2\lambda_{\text{comm}}(\Delta + \tau)N} (1 - e^{-\lambda_{\text{comm}}\tau}).$$

It can be shown that this also constitutes a lower bound for the other possible combinations of initial states: robot i is *waiting* and robot j is *moving*; robots i and j are both *moving*; and robots i and j are both *waiting*. ■

APPENDIX C

In this appendix we provide the proof of Proposition 2.8 which states that any pairwise-optimal partition is also a centroidal Voronoi partition.

Proof of Proposition 2.8: To create a contradiction, assume that $P \in \text{Part}_N(Q)$ is a pairwise-optimal partition but not a centroidal Voronoi partition. In other words, there exist components P_i and P_j in P and an element x of one component, say $x \in P_i$, such that

$$d_G(x, \text{Cd}(P_i)) > d_G(x, \text{Cd}(P_j)). \quad (5)$$

Choose P_j such that for all $k \neq j$

$$d_G(x, \text{Cd}(P_k)) \geq d_G(x, \text{Cd}(P_j)). \quad (6)$$

Let $s_{a,b}^G$ be a shortest path in G connecting a to b and let $m \in s_{x, \text{Cd}(P_j)}^G$ be the first element of the path starting from $\text{Cd}(P_j)$ which is not in P_j . Let ℓ be such that $m \in P_\ell$.

If $m = x$, then from (5) and the definition of $s_{x, \text{Cd}(P_j)}^G$ we have that

$$\begin{aligned} d_{P_i}(x, \text{Cd}(P_i)) &\geq d_G(x, \text{Cd}(P_i)) \\ &> d_G(x, \text{Cd}(P_i)) = d_{P_i \cup P_j}(x, \text{Cd}(P_j)) \end{aligned}$$

which, since $x \in P_i$, creates a contradiction of the fact that P is pairwise-optimal.

If $m \neq x$, then, given (6), one of these two conditions holds:

- (i) $d_G(m, \text{Cd}(P_\ell)) > d_G(m, \text{Cd}(P_j))$, or
- (ii) $d_G(m, \text{Cd}(P_\ell)) = d_G(m, \text{Cd}(P_j))$.

In the first case, we again have a contradiction using the same logic above with m in place of x . In the second case, we must further consider whether there exists a $s_{m, \text{Cd}(P_\ell)}^G$ such that every vertex in $s_{m, \text{Cd}(P_\ell)}^G$ is also in P_ℓ . If there is not such a path, then

$$d_{P_\ell}(m, \text{Cd}(P_\ell)) > d_G(m, \text{Cd}(P_\ell)) = d_{P_\ell \cup P_j}(m, \text{Cd}(P_j))$$

and we again have a contradiction as above. If there is such a path, then we can instead repeat this analysis using ℓ in place of j and considering the path formed by this $s_{m, \text{Cd}(P_\ell)}^G$ and the vertices in $s_{x, \text{Cd}(P_j)}^G$ after m . Since the next vertex playing the role of m must be closer to x , we will eventually find a vertex which creates a contradiction. ■

REFERENCES

- [1] R. Smith, J. Das, H. Heidarsson, A. Pereira, F. Arrichiello, I. Cetnic, L. Darjany, M.-E. Garneau, M. Howard, C. Oberg, M. Ragan, E. Seubert, E. Smith, B. Stauffer, A. Schnetzer, G. Toro-Farmer, D. Caron, B. Jones, and G. Sukhatme, "USC CINAPS builds bridges," *IEEE Robotics & Automation Magazine*, vol. 17, no. 1, pp. 20–30, 2010.
- [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] 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.
- [4] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith, "Dynamic vehicle routing for robotic systems," *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1482–1504, 2011.
- [5] A. K. Jain, M. N. Murty, and P. J. Flynn, "Data clustering: A review," *ACM Computing Surveys*, vol. 31, no. 3, pp. 264–323, 1999.
- [6] V. V. Vazirani, *Approximation Algorithms*. Springer, 2001.
- [7] P. O. Fjällström, "Algorithms for graph partitioning: A survey," *Linköping Electronic Articles in Computer and Information Science*, vol. 3, no. 10, 1998.
- [8] F. R. Adler and D. M. Gordon, "Optimization, conflict, and nonoverlapping foraging ranges in ants," *American Naturalist*, vol. 162, no. 5, pp. 529–543, 2003.
- [9] 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>.
- [10] S. P. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982. Presented at the 1957 Institute for Mathematical Statistics Meeting.
- [11] J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] R. Cortez, H. Tanner, and R. Lumia, "Distributed robotic radiation mapping," in *Experimental Robotics* (O. Khatib, V. Kumar, and G. Pappas, eds.), vol. 54 of *Springer Tracts in Advanced Robotics*, pp. 147–156, Springer, 2009.
- [16] 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.
- [17] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, 2006.
- [18] F. Bullo, R. Carli, and P. Frasca, "Gossip coverage control for robotic networks: Dynamical systems on the the space of partitions," *SIAM Journal on Control and Optimization*, Aug. 2010. Submitted. Available at <http://motion.me.ucsb.edu/pdf/2008u-bcf.pdf>.
- [19] 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, USA), pp. 225–232, Oct. 2009.
- [20] J. W. Durham, R. Carli, and F. Bullo, "Pairwise optimal coverage control for robotic networks in discretized environments," in *IEEE Conf. on Decision and Control*, (Atlanta, GA, USA), pp. 7286–7291, Dec. 2010.
- [21] H. Minc, *Nonnegative Matrices*. Wiley, 1988.
- [22] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry: Algorithms and Applications*. Springer, 2 ed., 2000.
- [23] B. Gerkey, R. T. Vaughan, and A. Howard, "The Player/Stage Project: Tools for multi-robot and distributed sensor systems," in *Int. Conference on Advanced Robotics*, (Coimbra, Portugal), pp. 317–323, June 2003.
- [24] J. G. Siek, L.-Q. Lee, and A. Lumsdaine, "Boost Graph Library," <http://www.boost.org>, July 2007. Version 1.34.1.
- [25] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, "Robust Monte Carlo localization for mobile robots," *Artificial Intelligence*, vol. 128, no. 1–2, pp. 99–141, 2001.
- [26] J. W. Durham and F. Bullo, "Smooth nearness-diagram navigation," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, (Nice, France), pp. 690–695, Sept. 2008.
- [27] R. Tempo, G. Calafiore, and F. Dabbene, *Randomized Algorithms for Analysis and Control of Uncertain Systems*. Springer, 2005.



Joseph W. Durham (S'10) received a B.A. magna cum laude in physics from Carleton College, Northfield, Minnesota, in 2004, and a M.Sc. and Ph.D. in control and dynamical systems from the University of California, Santa Barbara, in 2007 and 2011, respectively. In 2004 he worked on the Stanford University entry to the second DARPA Grand Challenge. In 2006-07 he was a teaching fellow with the L.E.A.P.S. program at Santa Barbara Junior High. He is currently a Research Scientist at Kiva Systems in North Reading, MA. His research interests include robotic coordination and navigation, resource allocation, and botting methods.



Ruggero Carli received the Laurea Degree in 2004 and the Ph.D. in 2007 in information engineering from the University of Padova, Padova, Italy. From 2008 through 2010 he was Post Doctoral Fellow with the Mechanical Engineering Department at the University of California in Santa Barbara. He is currently an Assistant Professor with the Department of Information Engineering, University of Padova. His research interests include control theory and, in particular, control under communication constraints, cooperative control and distributed estimation.



Paolo Frasca received the Ph.D. degree in applied mathematics from the Politecnico di Torino, Italy, in 2009. From 2008 to 2010, he was at the Università di Salerno and at the Istituto per le Applicazioni del Calcolo, C.N.R., Italy. He is currently with the Dipartimento di Matematica, Politecnico di Torino, Italy. His research interests include mathematical control theory and applications of probability and graph theory to distributed and networked systems.



Francesco Bullo (S'95-M'99-SM'03-F'10) received the Laurea degree summa cum laude in electrical engineering from the University of Padova, Padova, Italy, in 1994 and the Ph.D. degree in control and dynamical systems from the California Institute of Technology, Pasadena, in 1999.

From 1998 to 2004, he was an Assistant Professor with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana. He is currently a Professor with the Mechanical Engineering Department, University of California in Santa Barbara. His main research interests are multiagent networks with application to robotic coordination, distributed computing and power networks. He is the coauthor of *Geometric Control of Mechanical Systems* (Springer, 2004, ISBN: 0-387- 22195-6) and *Distributed Control of Robotic Networks* (Princeton, 2009, ISBN: 978-0-691-14195-4). He has published more than 200 papers in international journals, books, and refereed conferences.

Dr. Bullo has received the ONR Young Investigator Award, the Outstanding Paper Award for the IEEE Control Systems Magazine, and the AACC Hugo Schuck Best Paper Award. His students papers were finalists for the Best Student Paper Award at the IEEE Conference on Decision and Control (2002, 2005, 2007) and the American Control Conference (2005, 2006, 2010).