

Dynamic Partitioning and Coverage Control with Asynchronous One-to-Base-Station Communication

Joseph W. Durham

Ruggero Carli

Paolo Frasca

Francesco Bullo

Abstract— We propose algorithms to automatically deploy a group of mobile robots to provide coverage of a non-convex environment with communication limitations. In settings such as hilly terrain or for underwater ocean gliders, peer-to-peer communication can be impossible and frequent communication to a central base station may be impractical. This paper instead explores how to perform coverage control when each robot has only asynchronous and sporadic communication with a base station or, alternatively, with the rest of the team. Our approach evolves overlapping territories and provably converges to a centroidal Voronoi partition at equilibrium. We also describe how the use of overlapping territories allows our algorithm to smoothly handle dynamic changes to the robot team.

I. INTRODUCTION

In applications such as environmental monitoring [1] or warehouse logistics [2] a team of robots is asked to perform tasks over a large space. The distributed *environment partitioning problem* consists of designing control and communication laws for individual robots such that the team divides a space into regions in order to optimize the quality of service provided. *Coverage control* additionally optimizes the positioning of robots inside of a region.

Many existing coverage control algorithms assume that robots can communicate peer-to-peer [3], [4], but in some environments this is impractical. For example, underwater acoustic communication between ocean gliders is very low bandwidth and hilly or urban terrain can block radio communication. Instead, we present a coverage control algorithm for a team of robots who each have occasional contact with a central base station. This *one-to-base-station* communication model can represent ocean gliders surfacing to communicate with a tower [5], UAV data mules that periodically visit ground robots [6], or cost-mindful use of satellite or cellular communication. In addition, our algorithm optimizes the response time of the team to service requests in a non-convex environment represented by a graph, with optimality defined by a relevant “multi-center” cost function for overlapping territories. While the algorithm is given for one-to-base-station communication, it also works if each robot can occasionally broadcast a message directly to the whole team.

A broad discussion of partitioning and coverage control is presented in [7] which builds on the classic work of

Lloyd [8] on algorithms for optimal quantizer design through “centering and partitioning.” The Lloyd-type approach was first adapted for distributed coverage control in [3] and has since seen many variations, including non-convex environments [9] and “gossip” peer-to-peer communication [10]. The discretized non-convex domain considered here also appeared in [4] which looked at iterative optimal 2-partitioning.

Coverage control and territory partitioning have applications in many fields. In cyber-physical systems, applications include automated environmental monitoring [1], fetching and delivery [2], and other vehicle routing scenarios [11]. Coverage of discrete sets is closely related to the literature on data clustering and k -means [12], as well as the facility location or k -center problem [13].

There are three main contributions of this paper. First, we present the first coverage control algorithm for an asynchronous one-to-base-station communication model. This model is realistic and relevant for a variety of application domains, and the time delay between when robots communicate with the base station requires overlapping regions instead of a partition. Second, we prove that the algorithm converges to a centroidal Voronoi partition in finite time. Our Lyapunov argument is based on an adaptation of the standard partition-based coverage cost function. Overlapping regions also dictate changes to when to perform the classic Lloyd steps of centering versus territory exchange. Third, we describe how the algorithm can seamlessly handle the unscheduled arrival or departure of robots from the team. This feature leverages overlapping regions, and also eases integration of coverage control with task servicing.

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

In this Section we translate concepts used in partitioning of continuous environments to coverings on graphs. The one-to-base-station communication model in this paper requires overlapping coverings, instead of a partition.

A. Graph Distances

Let finite set Q be a set of points in a continuous environment. These points can represent small areas 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

This work was supported in part by NSF grants IIS-0904501 and CPS-1035917.

Joseph W. Durham and Francesco Bullo are with the Department of Mechanical Engineering, University of California, Santa Barbara, CA, 93106 (joey, 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, Italy paolo.frasca@polito.it

$w_e > 0$ be the weight of edge e . We assume that $G(Q)$ is connected and think of the edge weights as travel distances between nearby points.

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 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. Coverings of Graphs

We will be covering Q with n subsets or regions which will each be owned by an individual agent.

Definition II.1 (*n*-Covering) Given the graph $G(Q) = (Q, E, w)$, we define a *n*-covering of Q as a collection $P = \{P_i\}_{i=1}^n$ of subsets of Q such that:

- (i) $\bigcup_{i=1}^n P_i = Q$;
- (ii) $P_i \neq \emptyset$ for all $i \in \{1, \dots, n\}$;

Let $\text{Cov}_n(Q)$ to be the set of *n*-coverings of Q .

Note that a vertex in Q may belong to multiple subsets in P , i.e., a vertex may be covered by multiple agents. This fact is an important change from prior work [4].

We also have use for the concept of a partition of Q .

Definition II.2 (*n*-Partition) A *n*-partition is a *n*-covering with the additional property that:

- (iii) if $i \neq j$, then $P_i \cap P_j = \emptyset$.

Let $\text{Part}_n(Q)$ to be the set of *n*-partitions of Q .

Among the ways of covering Q , there is one which is 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 $d_G(k, c_i) \leq d_G(k, c_j)$, $\forall j \neq i$. The elements of c are said to be the generators of the Voronoi partition. Note that the Voronoi partition generated by c is not unique since how to assign tied vertices is unspecified.

C. Cost Functions

Let *weight function* $\phi : Q \rightarrow \mathbb{R}_{>0}$ be a bounded positive function which assigns a relative weight to each element of Q . The *one-center function* \mathcal{H}_1 gives the cost for a robot to cover a subset $A \subset Q$ from a vertex $h \in A$ with relative prioritization set by ϕ :

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

A technical assumption is needed to define the *generalized centroid* of a subset. We assume from now on that a *total order* relation, $<$, is defined on Q : hence, we can denote $Q = \{1, \dots, |Q|\}$. With this assumption we can deterministically pick a centroid in P_i which minimizes \mathcal{H}_1 as follows.

Definition II.3 (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}_1 , i.e.,

$$C(A) := \underset{h \in A}{\operatorname{argmin}} \mathcal{H}_1(h; A).$$

Furthermore, we define the map $\text{Cd} : 2^Q \rightarrow Q$ such that $\text{Cd}(A) := \min\{c \in C(A)\}$. 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 set always belongs to the set. With a slight notational abuse, we define $\text{Cd} : \text{Cov}_n(Q) \rightarrow Q^n$ as the map which associates to a covering the vector of the centroids of its elements.

With these notions we can define the multi-center function $\mathcal{H}_{\max} : Q^n \times \text{Cov}_n(Q) \rightarrow \mathbb{R}_{\geq 0}$ to measure the cost for n robots to cover a *n*-covering P from the vertex set $c \in Q^n$:

$$\mathcal{H}_{\max}(c, P) = \frac{1}{|Q|} \sum_{k \in Q} \max_i \{d_G(c_i, k) \mid k \in P_i\} \phi(k)$$

We aim to minimize the performance function \mathcal{H}_{\max} with respect to both the covering P and the vertices c . In the motivational scenario we are considering, each robot will periodically be asked to perform a task somewhere in its region with tasks located according to distribution ϕ . When idle, the robots would position themselves at the vertices c . By minimizing \mathcal{H}_{\max} , the robot team would minimize the expected distance between a task and the furthest robot which can service the task.

Proposition II.4 (Properties of \mathcal{H}_{\max}) Let $P \in \text{Cov}_n(Q)$, $P' \in \text{Part}_n(Q)$, and $c \in Q^n$ such that $c_i \in P'_i \subseteq P_i \forall i$. Let $c' \in Q^n$ such that $c'_i \in C(P'_i) \forall i$. Then the following statements hold:

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

The second inequality is strict if any $c_i \notin C(P'_i)$.

Proof: The first statement is a straightforward consequence of the restriction that $P'_i \subseteq P_i$ and that \mathcal{H}_{\max} uses the maximum cost over i . The second statement is a result of the fact that, since P' is a partition, $\mathcal{H}_{\max}(c, P') = \frac{1}{|Q|} \sum_i \mathcal{H}_1(c_i; P'_i)$. ■

Proposition II.4 motivates the following definition.

Definition II.5 (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 $c_i \in C(P_i) \forall i$.

For a given environment Q , a pair made of a centroidal Voronoi partition and the corresponding vector of centroids is locally optimal in the following sense: \mathcal{H}_{\max} cannot be reduced by changing either P or c independently. Therefore, if the team of robots position themselves at the centroids of a centroidal Voronoi partition, then they (locally) optimize their coverage of Q as measured by \mathcal{H}_{\max} .

III. MODEL, PROBLEM, AND PROPOSED SOLUTION

A. One-to-Base-Station Robotic Network Model

We are given a team of n robotic agents and a central base station. Each agent $i \in \{1, \dots, n\}$ is required to have the following basic computation capabilities:

- (C1) agent i can identify itself to the base station; and
- (C2) agent i has a processor with the ability to store $S_i \subset G(Q)$ and a center $s_i \in S_i$.

Each agent $i \in \{1, \dots, n\}$ is assumed to communicate with the base station according to the *asynchronous one-to-base-station communication model* described as follows:

- (C3) there exists a finite upper bound Δ on the time between communications between i and the base station. For simplicity, we assume no two agents communicate with the base station at the same time.

The base station must have the following capabilities:

- (C4) it can store an arbitrary n -covering of Q , $P = \{P_i\}_{i=1}^n$ and a list of centroids $c \in Q^n$; and
- (C5) it can perform computations on subgraphs of $G(Q)$.

B. Problem Statement

Assume that, for all $t \in \mathbb{R}_{\geq 0}$, each agent $i \in \{1, \dots, n\}$ maintains in memory a subset $S_i(t)$ of environment Q and a vertex $s_i(t) \in S_i(t)$. Our goal is to iteratively update the covering $S(t) = \{S_i(t)\}_{i=1}^n$ and the centers $s(t) = \{s_i(t)\}_{i=1}^n$ while solving the optimization problem:

$$\min_{s \in Q^n} \min_{S \in \text{Cov}_n(Q)} \mathcal{H}_{\max}(s, S), \quad (1)$$

subject to the constraints imposed by the robot network model with asynchronous one-to-base-station communication from Section III-A.

C. The One-to-Base Coverage Algorithm

To solve the minimization problem (1), we introduce the following *One-to-Base Coverage Algorithm*.

One-to-Base Coverage Algorithm

The base station maintains in memory an n -covering $P = \{P_i\}_{i=1}^n$ and a vector $c = (c_i)_{i=1}^n$, while each robot maintains in memory a set S_i and a vertex s_i . At $t = 0$, let $P(0) \in \text{Cov}_n(Q)$, $S(0) = P(0)$, and let all $c_i(0)$'s be distinct. Assume that at time $t \in \mathbb{R}_{>0}$, robot i communicates with the base station. Let P^+ , c^+ , S_i^+ , and s_i^+ be the values after the communication. Then the base station executes the following actions while communicating with i :

- 1: **if** $\mathcal{H}_1(\text{Cd}(P_i); P_i) < \mathcal{H}_1(c_i; P_i)$ and $\text{Cd}(P_i) \neq c_j$ for every $j \neq i$ **then**
 - 2: update $c_i^+ := \text{Cd}(P_i)$
 - 3: **else**
 - 4: $c_i^+ := c_i$
 - 5: tell agent i to set $S_i^+ := P_i$ and $s_i^+ := c_i^+$
 - 6: **for** each agent $j \neq i$ **do**
 - 7: compute the sets

$$P_{i \rightarrow j} := \{x \in P_i : d_G(x, c_j) < d_G(x, c_i^+)\}$$

$$P_{j \rightarrow i}^* := \{x \in P_j \cap P_i : d_G(x, c_j) \geq d_G(x, c_i^+)\}$$
 - 8: $P_j^+ := (P_j \setminus P_{j \rightarrow i}^*) \cup P_{i \rightarrow j}$
-

Observe that $P_{i \rightarrow j}$ contains the cells of P_i which are closer to c_j , whereas $P_{j \rightarrow i}^*$ contains only the cells in both P_i and P_j which are either closer to c_i^+ or tied. Also, only the centroid of robot i is updated. Finally, note that the algorithm is independent of robot positions, so the robots are free to move or perform tasks in their regions.

Remark III.1 *The One-to-Base algorithm can be adapted to the scenario where each robot can occasionally broadcast a message to the team. Robot i would update its centroid and broadcast s_i^+ and S_i , then every other robot j would update S_j following lines 7 and 8 above. Those robots for which $S_i \cup S_j$ is connected must receive the broadcast for the convergence property to hold, the others are not required.*

D. Convergence Property

In this subsection we characterize the convergence of the One-to-Base Coverage Algorithm.

Theorem III.2 (Convergence Property) *Consider a network consisting of n robots endowed with computation capacities (C1), (C2) and communication capacity (C3), and a base station with capacities (C4) and (C5). Assume the network implements the One-to-Base Coverage Algorithm. Then the resulting evolution $(s, S) : \mathbb{R}_{\geq 0} \rightarrow Q^n \times \text{Cov}_n(Q)$ converges in finite time to a pair (s^*, S^*) composed of a centroidal Voronoi partition S^* generated by s^* .*

Remark III.3 *The fact that at least one centroidal Voronoi partition exists for any graph is an additional consequence of Theorem III.2.*

IV. CONVERGENCE PROOFS

This section is devoted to proving Theorem III.2. The convergence proof is based on applying Lemma A.1 to the evolution given by the One-to-Base Coverage Algorithm. To do so, we must describe the algorithm using a set valued-map and find a Lyapunov function.

A. Set-valued Map

With the definitions of a set of centroids and of the One-to-Base Coverage Algorithm, we have that the algorithm is well-posed in the following sense.

Proposition IV.1 (Well-posedness) *Let $P \in \text{Cov}_n(Q)$ and $c \in Q^n$ such that $c_i \in P_i$ and $c_i \neq c_j$ for all i and all $j \neq i$. Then, P^+ and c^+ produced by the One-to-Base Coverage Algorithm meet the same criteria.*

With this result, we can state the One-to-Base Coverage Algorithm as a set valued map. For any $i \in \{1, \dots, n\}$, we define the map $T_i : Q^n \times \text{Cov}_n(Q) \rightarrow Q^n \times \text{Cov}_n(Q)$ by

$$T_i(c, P) = \{ \{c_1, \dots, c_i^+, \dots, c_n\}, \{P_1^+, \dots, P_i, \dots, P_n^+\} \},$$

where c_i^+ and P^+ are defined per the algorithm when i is the communicating robot. Then, we can define the set-valued map $T : Q^n \times \text{Cov}_n(Q) \rightrightarrows Q^n \times \text{Cov}_n(Q)$ by

$$T(c, P) = \{T_1(c, P), \dots, T_n(c, P)\}.$$

Thus, the dynamical system defined by the application of the algorithm is described by $\{c^+, P^+\} \in T(c, P)$.

B. Lyapunov Function

For our Lyapunov argument we need the following definitions. Let $M(P)$ be the set of vertices which are owned by multiple agents. Let \mathcal{H}_{\min} be a cost function defined similarly to \mathcal{H}_{\max} but sum minimum coverage costs over all agents:

$$\mathcal{H}_{\min}(c, P) = \frac{1}{|Q|} \sum_{k \in Q} \min_i \{d_G(c_i, k) \mid k \in P_i\} \phi(k).$$

Proposition IV.2 (Decreasing Functions) *Let P be a n -covering of Q and c be a set of centroids for P . Let $(c^+, P^+) \in T(c, P)$. If $c^+ \neq c$ or $P^+ \neq P$, then one of these conditions holds:*

- (i) $\mathcal{H}_{\max}(c^+, P^+) < \mathcal{H}_{\max}(c, P)$;
- (ii) $\mathcal{H}_{\max}(c^+, P^+) = \mathcal{H}_{\max}(c, P)$ and $\mathcal{H}_{\min}(c^+, P^+) < \mathcal{H}_{\min}(c, P)$; or
- (iii) $\mathcal{H}_{\max}(c^+, P^+) = \mathcal{H}_{\max}(c, P)$, $\mathcal{H}_{\min}(c^+, P^+) = \mathcal{H}_{\min}(c, P)$, and $|M(P^+)| < |M(P)|$.

Proof: Consider the situation where there are just two agents i and j . Without loss of generality, assume that agent i contacts the base station at time t .

We start with the case where $c_i^+ = c_i$. First, consider when the change to P includes the addition of cells in $P_{i \rightarrow j}$ to P_j . Such a change necessarily decreases \mathcal{H}_{\min} while \mathcal{H}_{\max} is unchanged. Next, if the change to P occurs because of the removal of cells in $P_{j \rightarrow i}^*$ from P_j , then \mathcal{H}_{\max} does not increase, \mathcal{H}_{\min} is unchanged, and $|M|$ necessarily decreases.

Next, we show that if $c_i^+ \neq c_i$, then $\mathcal{H}_{\max}(c^+, P^+) < \mathcal{H}_{\max}(c, P)$. First, given a $P \in \text{Cov}_n(Q)$, let $P_{\max} = \{P_{i, \max}\}_{i=1}^n$ be a partition of Q such that for all i :

$$P_{i, \max} = \left\{ v \in P_i \mid \begin{array}{l} v \notin P_j \ \forall j \neq i, \text{ or} \\ i = \min \{ \text{argmax}_j \{ d_G(c_j, v) \mid v \in P_j \} \} \end{array} \right\}.$$

Note that $P_{i, \max}$ is a function of P_i, P_j, c_i , and c_j .

With the P_{\max} definition, we can rewrite \mathcal{H}_{\max} as:

$$\mathcal{H}_{\max}(c, P) = \frac{1}{|Q|} \sum_i \mathcal{H}_1(c_i, P_{i, \max})$$

Using this new form, the initial cost to cover Q by i and j is given by (ignoring $\frac{1}{|Q|}$ for simplicity):

$$\begin{aligned} \mathcal{H}_{\max}(c, P) = & \mathcal{H}_1(c_i, P_{i, \max}) + \mathcal{H}_1(c_j, P_{j, \max} \setminus P_i) \\ & + \mathcal{H}_1(c_j, P_{j, \max} \cap P_i). \end{aligned}$$

During the update c_i and P_j change, meaning that:

$$\begin{aligned} \mathcal{H}_{\max}(c^+, P^+) = & \mathcal{H}_1(c_i^+, P_{i, \max}^+) + \mathcal{H}_1(c_j, P_{j, \max}^+ \setminus P_i) \\ & + \mathcal{H}_1(c_j, P_{j, \max}^+ \cap P_i). \end{aligned}$$

The algorithm T ensures that if $c_i^+ \neq c_i$, then:

$$\mathcal{H}_1(c_i^+, P_i) < \mathcal{H}_1(c_i, P_i). \quad (2)$$

However, it is possible that the relevant cost for i has increased, i.e., that $\mathcal{H}_1(c_i^+, P_{i, \max}^+) > \mathcal{H}_1(c_i, P_{i, \max})$. We

will show that any such increase is necessarily smaller in magnitude than the decrease in the cost to cover for j .

Two observations: First, $P_{j, \max}^+ \cap P_i = \emptyset$ by how we choose $P_{j, \max}^+$, meaning that $\mathcal{H}_1(c_j, P_{j, \max}^+ \cap P_i)$ is zero. Second, the set of vertices owned by j but not by i has not changed, meaning that $\mathcal{H}_1(c_j, P_{j, \max}^+ \setminus P_i) = \mathcal{H}_1(c_j, P_{j, \max} \setminus P_i)$. This leaves us wanting to show that:

$$\mathcal{H}_1(c_i^+, P_{i, \max}^+) < \mathcal{H}_1(c_i, P_{i, \max}) + \mathcal{H}_1(c_j, P_{j, \max} \cap P_i).$$

We can write set P_i as:

$$\begin{aligned} P_i &= P_{i, \max} \cup (P_{j, \max} \cap P_i) \\ &= P_{i, \max}^+ \cup (P_{j, \max}^+ \cap P_i) = P_{i, \max}^+. \end{aligned}$$

Using these equivalences, we can rewrite (2) as:

$$\mathcal{H}_1(c_i^+, P_{i, \max}^+) < \mathcal{H}_1(c_i, P_{i, \max}) + \mathcal{H}_1(c_j, P_{j, \max} \cap P_i).$$

Then, using the definition of $P_{j, \max}$ we conclude that:

$$\begin{aligned} \mathcal{H}_1(c_i^+, P_{i, \max}^+) &< \mathcal{H}_1(c_i, P_{i, \max}) + \mathcal{H}_1(c_j, P_{j, \max} \cap P_i) \\ &< \mathcal{H}_1(c_i, P_{i, \max}) + \mathcal{H}_1(c_j, P_{j, \max} \cap P_i). \end{aligned}$$

Nothing in this analysis is exclusive to the two agent scenario. Following the same logic, it can be shown that:

$$\mathcal{H}_1(c_i^+, P_{i, \max}^+) - \mathcal{H}_1(c_i, P_{i, \max}) < \sum_{j \neq i} \mathcal{H}_1(c_j, P_{j, \max} \cap P_i),$$

meaning that any increase in the cost to cover for agent i from a centroid update is more than offset by decreases to the cost to cover from the territory updates of those agents who owned cells in P_i . ■

We can form a Lyapunov function using Proposition IV.2 as follows. Since Q is a finite set, there exist only a finite number of possible values for \mathcal{H}_{\max} , \mathcal{H}_{\min} , and $|M|$. Let ϵ_x and ϵ_n be the magnitude of the smallest possible difference between two values of \mathcal{H}_{\max} and \mathcal{H}_{\min} , respectively. Let α_n and α_M be larger than twice the maximum possible values of \mathcal{H}_{\min} and $|M|$, respectively. Consider the following function $U : Q^n \times \text{Cov}_n(Q) \rightarrow \mathbb{R}_{\geq 0}$:

$$U(c, P) = \mathcal{H}_{\max}(c, P) + \frac{\epsilon_x}{\alpha_n} \mathcal{H}_{\min}(c, P) + \frac{\epsilon_x}{\alpha_n} \frac{\epsilon_n}{\alpha_M} |M(P)|.$$

With this scaling of \mathcal{H}_{\min} and $|M|$, when \mathcal{H}_{\max} decreases then U necessarily also decreases, and similarly if \mathcal{H}_{\max} is constant but \mathcal{H}_{\min} decreases. We further have the following bound on changes to U .

Proposition IV.3 (Lyapunov Function) *Let $(c', P') \in T(c, P)$. Then, either $(c', P') = (c, P)$ or $U(c', P') \leq U(c, P) - \frac{\epsilon_x \epsilon_n}{\alpha_n \alpha_M}$.*

C. Characterization of Fixed Points

One consequence of Proposition IV.3 is that the maps T_i have at least one common fixed point. The following Proposition characterizes the fixed points for $T(c, P)$, defined¹ as the pairs (c, P) where $\{(c, P)\} = T(c, P)$ or, equivalently, as the pairs which are a fixed point of every map T_i .

¹The standard definition of fixed point for a set-valued map (which we do not use in this paper) consists in the weaker condition $(c, P) \in T(c, P)$.

Proposition IV.4 (Fixed Points) *Let $(c, P) \in Q^n \times \text{Cov}_n(Q)$ be a fixed point of T . Then, P is a centroidal Voronoi partition of Q generated by c . Moreover, every such centroidal Voronoi partition is a fixed point for T .*

Proof: If P is not a partition, then $P_{j \rightarrow i}^* \neq \emptyset$ for some $i \neq j$. If P is a partition but not a Voronoi partition generated by c , then $P_{i \rightarrow j} \neq \emptyset$ for some $i \neq j$. Finally, if P is a Voronoi partition generated by c but $c_i \notin C(P_i)$ for any i , then c_i will change when i communicates with the base station.

Next, we show that every centroidal Voronoi partition is a fixed point. If $c_i \in C(P_i)$ for all i , then $c_i^+ = c_i$ for all T_i . If P is a Voronoi partition generated by c , then $P_{i \rightarrow j} = \emptyset$, $P_{j \rightarrow i}^* = \emptyset$, and thus $P^+ = P$ for all T_i . ■

D. Convergence of $P(t)$

The proof continues with the application of Lemma A.1 in Appendix A to $(c(t), P(t))$. Since the algorithm $T : Q^n \times \text{Cov}_n(Q) \rightrightarrows Q^n \times \text{Cov}_n(Q)$ is well-posed, we have that $Q^n \times \text{Cov}_n(Q)$ is strongly positively invariant. This fact implies that assumption (i) of Lemma A.1 is satisfied. Invoking Proposition IV.3, we conclude that $U(c, P)$ fulfills assumption (ii). Finally, the communication model (C3) assures that assumption (iii) is met.

Hence, we are in the position to apply Lemma A.1 and conclude the following result.

Proposition IV.5 (Convergence of $P(t)$) *The evolution of the One-to-Base Coverage Algorithm $(c(t), P(t))$, generated by the map T , converges in finite time to the intersection of the equilibria of the maps T_i , which is the set of pairs (c, P) where P is a centroidal Voronoi partition generated by c . In particular, $P(t)$ converges in finite time to one centroidal Voronoi partition.*

E. Convergence of Robot Covering

So far we have discussed the properties of the covering P held by the base station. Here we extend these arguments to the covering S held by the robots. First, we show that S is indeed a covering of Q .

Proposition IV.6 (Well-posedness of S) *Let S be a n -covering of Q . Then, S^+ produced by the One-to-Base Coverage Algorithm is also a n -covering.*

Proof: Let $s \in Q$. If there exist times $t_1 < t_2$ such that $q \in S_i(t_1)$ and $q \notin S_i(t_2)$, then there exists a $\tilde{t} \in [t_1, t_2)$ such that $q \notin P_i(\tilde{t}^+)$. By how the update of $P(t)$ is defined, this implies that some agent $j \neq i$ with $q \in P_j(\tilde{t})$ communicates to the base station at time \tilde{t} . But since $S_j(\tilde{t}^+) = P_j(\tilde{t})$, we have that $q \in S_j(\tilde{t}^+)$. Therefore, q must belong to some region of $S(t)$ for all t . ■

We are now ready to conclude our convergence proof. *Proof:* [of Theorem III.2]. The definition of the One-to-Base Coverage Algorithm implies that if there exists $\tau \in \mathbb{N}$ such that $P(t) = \bar{P} \in \text{Cov}_n(Q)$ for $t \geq \tau$, then $S(t) = \bar{P}$ for $t \geq \tau + \Delta$. As an immediate consequence of this fact, the convergence properties of $P(t)$, stated in Proposition IV.5, are inherited by $S(t)$. ■

Evolving overlapping coverings enables simple handling of dynamic arrivals, departures, and even the disappearance of robots. While departure or disappearance can increase \mathcal{H}_{\max} , such an increase is only a transient and, with the following additions, the system will converge to a centroidal Voronoi partition in finite steps after such an event.

Arrival: When a new robot i communicates with the base station, it can be assigned any initial P_i desired. Possibilities include adding all vertices within a set distance of its initial position or assigning it just the single vertex which has the highest coverage cost in Q .

Departure: A robot i might announce to the base station that it is departing, perhaps to recharge its batteries or to perform some other task. In this situation, the base station can simply add P_i to the territory of the next robot it talks to before executing the normal steps of the algorithm.

Disappearance: The disappearance or failure of a robot i can be detected if it does not communicate with the base station for longer than Δ . If this occurs, then the departure procedure above can be triggered. Should i reappear later, it can be handled as a new arrival or given its old territory.

VI. NUMERICAL RESULTS

To demonstrate the utility of the One-to-Base Coverage Algorithm, we implemented it using the open-source Player/Stage robot control system and the Boost Graph Library (BGL). All results presented here were generated using Player 2.1.1, Stage 2.1.1, and BGL 1.34.1.

One illustrative example is shown in Figure 1. The environment contains three obstacles drawn in black and four robots tasked with providing coverage of the free space around the obstacles. This free space is modeled using an occupancy grid with a $0.6m$ resolution which was chosen so that the robots could fit inside of a grid cell. The grid is converted into a graph by making each free cell a vertex and connecting edges between cells which border each other. To compute distances in this uniform edge weight graph we extended the BGL breadth-first search routine with a distance recorder event visitor.

For this example we chose a random robot to communicate with the base station at each iteration, while ensuring that no robot went unselected for more than 8 rounds. In the covering shown in the second panel of Figure 1, the light blue robot on the top left and the dark blue robot on the middle left both own some vertices also claimed by the circled orange robot. The third panel shows the result after the orange robot communicates with the base station: the orange robot's centroid has been updated and both blue robots have relinquished their claim to vertices closer to orange.

The final centroidal Voronoi partition in the fourth panel is reached after 25 iterations. The final coverage cost was $1.82m$, an improvement of 59%. Since each robot initially covers the entire environment, this also represents the improvement from using four robots instead of one to provide coverage in this environment.




Fig. 1. Simulation of four robots partitioning an environment with three black obstacles. The free space of the environment is modeled using the indicated occupancy grid where each cell is a vertex in the resulting graph. On the left, each robot starts owning the entire environment and positioned at its initial unique centroid. The middle frames show an intermediate state of the covering P and the result of an update when the circled robot contacts the base station. The centroids are marked with an X and the boundary of each robot's territory drawn in its color. Some cells are on the boundary of multiple territories and for these we draw superimposed robot colors. The final partition is shown at right.

VII. DISCUSSION & CONCLUSION

We have described the One-to-Base Coverage Algorithm which can drive territory ownership among a team of robots in a non-convex environment to a centroidal Voronoi partition in finite time given only occasional contact between each robot and a central base station. Here we have focused on dividing territory, but the algorithm can easily be combined with methods to provide a service over Q , as in [11].

In practical use, between the times that a robot communicates with the base station it could take sample measurements, pick up packages, or perform other tasks. When a robot communicates to the base station, it could transmit any information it has gathered about the environment and then receive its updated territory and a list of tasks to perform. When idle, a robot would position itself at the centroid of its territory. If tasks appear according to the distribution ϕ (which could evolve over time), then by minimizing cost function \mathcal{H}_{\max} the algorithm also minimizes the expected distance between a task and the furthest idle robot which might be assigned the task.

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 [10, Theorem 4.3].

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

Lemma A.1 (Convergence under persistent switches)

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 $T(x) = \{T_1(x), \dots, T_m(x)\}$ and let $\{x_n\}_{n \in \mathbb{Z}_{\geq 0}}$ be an evolution of T . Assume that:

- (i) there exists $W \subseteq X$ 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) for all $i \in \{1, \dots, m\}$, there exists an increasing sequence of times $\{n_k \mid k \in \mathbb{Z}_{\geq 0}\}$ such that $x_{n_{k+1}} = T_i(x_{n_k})$ and $(n_{k+1} - n_k)$ is bounded.

If $x_0 \in W$, there exists $c \in \mathbb{R}$ and $N \in \mathbb{N}$ such that for all $n \geq N$, the evolution $x_n = \bar{x}$ where \bar{x} belongs to the set $(F_1 \cap \dots \cap F_m)$, where $F_i = \{w \in W \mid T_i(w) = w\}$ is the set of fixed points of T_i in W , $i \in \{1, \dots, m\}$.

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] 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.
- [4] 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.
- [5] A. Pereira, H. Heidarsson, C. Oberg, D. Caron, B. Jones, and G. Sukhatme, "A communication framework for cost-effective operation of AUVs in coastal regions," in *Field and Service Robotics* (A. Howard, K. Iagnemma, and A. Kelly, eds.), vol. 62 of *Tracts in Advanced Robotics*, pp. 433–442, Springer, 2010.
- [6] R. C. Shah, S. Roy, S. Jain, and W. Brunette, "Data MULEs: modeling and analysis of a three-tier architecture for sparse sensor networks," *Ad Hoc Networks*, vol. 1, no. 2-3, pp. 215–233, 2003.
- [7] 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>.
- [8] 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.
- [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] F. Bullo, R. Carli, and P. Frasca, "Gossip coverage control for robotic networks: Dynamical systems on the space of partitions," *SIAM Journal on Control and Optimization*, Aug. 2010. Submitted. Available at <http://motion.me.ucsb.edu/pdf/2008u-bcf.pdf>.
- [11] 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.
- [12] 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.
- [13] V. V. Vazirani, *Approximation Algorithms*. Springer, 2001.