

Multiagent coverage algorithms with gossip communication: control systems on the space of partitions

Paolo Frasca

Ruggero Carli

Francesco Bullo

Abstract—Deployment, coverage and partitioning are fundamental tasks for robotic networks. Recently proposed algorithms achieve these tasks under a critical assumption: information is exchanged synchronously among all agents and long-range communication is possibly required. This work proposes novel deployment and partitioning algorithms that require only asynchronous pairwise (so-called gossip) communication. Which robot pair communicates at any given time may be selected deterministically or randomly. A key novel idea is the description of the coverage control problem as a control system on the space of partitions – in other words, we study the evolution of the regions assigned to each agent, rather than the evolution of the agents’ positions. The novel gossip algorithms are shown to converge to multicenter Voronoi partitions through various results of independent interest: we establish the compactness of the space of partitions, the continuity of certain geometric maps (e.g., the Voronoi and the centroid maps), and two convergence theorems for switching dynamical systems on metric spaces.

I. INTRODUCTION

This paper considers deployment and partitioning problems for robotic networks, that is, groups of robots that have basic motion, communication and computation capacities and that coordinate their actions based on simple interaction laws and protocols. The *deployment problem* for a robotic network amounts to the design of control and communication laws that lead the robots to be optimally placed in an environment of interest; the usual approach consists of identifying an appropriate “network cost function” that measures the deployment quality of a given configuration and designing control and communication laws that optimize this measure. The *partitioning problem* is the design of control and communication laws that lead the robots to optimally partition the environment into subregions of interest; even here the objective is usually achieved through the design of appropriate cost functions. Coverage control algorithms typically solve both deployment and partitioning problems simultaneously.

Broad discussions about distributed control algorithms for coverage, deployment and partitioning are presented in [1], [2]; these discussions build on the classic work by Lloyd [3] on “centering and partitioning” algorithms for optimal quantizer design. In [4] partitioning policies are shown to achieve optimal load balancing in vehicle routing

problems, i.e., problems in which a robotic network has the task of visiting points generated over time by a stochastic process. Multicenter Voronoi partitions are shown in [5] to be asymptotically optimal for estimation of stochastic spatial fields. Convergence to multicenter Voronoi partitions is established in [6] for a class of communication-less sensor-based algorithm (related to the classic clustering work [7]). Dynamic environments and corresponding dynamic coverage problems are treated in [8]. Nonconvex environments, maximization of detection probability and heterogeneous robotic networks are discussed in [9], [10], [11].

In our work we adopt methods from distinct disciplines. For example, the pairwise “gossip” approach to agents communication is widely adopted in the wireless communication and consensus literature; see [12] and subsequent works. Additionally, we adopt various tools from topology; the application of topological methods to multiagent systems and distributed coverage verification has received much recent attention [13], [14]. Finally, we consider control systems on a non-Euclidean state space; the interest for non-Euclidean spaces has a rich history in nonlinear control theory and robotics, e.g., see the early work [15].

The main contributions of the present paper are three. First, we describe coverage control algorithms in a novel way. Classically, the state space for the coverage algorithms are the agent positions: based on their positions, the agents apportion the environment into regions, which are assigned to each agent. In our approach, the agents positions are no longer the main concern: the state space is a space of partitions of the given environment. We discuss important properties of such a space, namely its compactness with respect to a suitable metric, and the continuity of several functions defined on it.

Second, as key motivating application, we devise a novel algorithm for coverage optimization, a “gossip” algorithm, in which only one pair of agents communicates per time step. We do this, because we know that reducing the communication burden is a critical issue for coverage control: indeed pairwise communication can be more effective in practical situations if connections between agents are not guaranteed to be fully reliable. Additionally, this asynchronous pairwise mechanism may be implemented also for robots with limited range communication. We propose a “novel random destination + wait” algorithm that achieves the required persistent communication requirements.

Third, we provide convergence theorems which extend the LaSalle invariance principle to a special class of set-valued maps on metric spaces. Convergence to a certain set of fixed points is achieved under uniform deterministic or stochastic persistency conditions. Applying these extensions

This work was supported in part by ARO MURI grant W911NF-05-1-0219, ONR grant N00014-07-1-0721 and NSF grant CNS-0834446.

Paolo Frasca was with the Dipartimento di Matematica, Politecnico di Torino, and is now with the DIIMA, Università di Salerno and with the Istituto per le Applicazioni del Calcolo-CNR, Roma, Italy, paolo.frasca@polito.it. Ruggero Carli and Francesco Bullo are with the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, Santa Barbara, CA 93106, USA, [{carlirug, bullo}@engineering.ucsb.edu"> {carlirug, bullo}@engineering.ucsb.edu](mailto).

of the LaSalle invariance principle and the properties of the state of partitions, we are able to give conditions for the proposed algorithm to converge to the critical points of a natural cost functional.

A. Organization and notations

The paper is structured as follows. In Section II we formally describe the coverage control problem. In Section III we present the gossip coverage algorithm, we state its convergence properties, and we show simulation results. Section IV contains the convergence theorems extending the LaSalle invariance principle; Section V describes the space of partitions; and Section VI states the continuity properties of the relevant maps and functions. Some conclusions are given in Section VII. In the interest of brevity we removed all proofs and placed them in a freely available technical report at <http://arXiv.org/abs/0903.3642>.

We let $\mathbb{R}_{>0}$ and $\mathbb{R}_{\geq 0}$ denote the set of positive and non-negative real numbers, respectively, and $\mathbb{Z}_{\geq 0}$ denote the set of non-negative integer numbers. Given a subset A of the Euclidean space \mathbb{R}^d , we let $\text{int}(A)$ denote its interior, \bar{A} denote its closure, ∂A denote its boundary and $\text{diam}(A)$ its diameter. Given two sets X and Y , a *set-valued map* $T : X \rightrightarrows Y$ associates to an element of X a subset of Y .

II. COVERAGE OPTIMIZATION AND DISTRIBUTED CONTROL VIA MULTICENTER FUNCTIONS

We are given a group of robots (also called agents) with limited communication and sensing capabilities, and an environment, and we want the agents to deploy in the area in an optimal way. The environment is apportioned into smaller regions, each assigned to an agent. Iteratively, the partition, and the agents configuration, are updated in a way to minimize a cost functional, which depends on the current partition and agents' positions.

A. Partitions, centroids and multicenter optimization

In what follows, let the environment to apportion be Q , a compact convex subset of \mathbb{R}^d with non-empty interior. Partitions of Q are defined as follows.

Definition II.1 (Partition) *An N -partition of Q , denoted by $v = \{v_i\}_{i=1}^N$, is a collection of N subsets of Q with the following properties:*

- (i) *each set v_i , $i \in \{1, \dots, N\}$, is closed, has non-empty interior, and its boundary has measure zero;*
- (ii) *$\text{int}(v_i) \cap \text{int}(v_j)$ is empty whenever $i \neq j$; and*
- (iii) *$\cup_{i \in \{1, \dots, N\}} v_i = Q$.*

We let \mathcal{V}_N denote the set of N -partitions of Q .

Let $p = (p_1, \dots, p_N) \in Q^N$ denote the position of N agents in the environment Q . Given $v \in \mathcal{V}_N$ and almost any $p \in Q^N$, each agent is naturally in one-to-one correspondence with an element of v ; specifically we sometimes refer to v_i as the *dominance region* of agent $i \in \{1, \dots, N\}$.

On Q , we define a *density function* to be a bounded measurable positive function $\phi : Q \rightarrow \mathbb{R}_{>0}$ and a *performance function* to be a locally Lipschitz, monotone increasing and convex function $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. With these notions, we

next define the *multicenter function* $\mathcal{H}_{\text{multicenter}} : \mathcal{V}_N \times Q^N \rightarrow \mathbb{R}_{\geq 0}$ by

$$\mathcal{H}_{\text{multicenter}}(v, p) = \sum_{i=1}^N \int_{v_i} f(\|p_i - q\|) \phi(q) dq. \quad (1)$$

We aim to minimize this function with respect to both the partition v and the locations p .

Remark II.2 (A word about locational optimization)

The function $\mathcal{H}_{\text{multicenter}}$ has the following interpretation. Given an agent at location p_i , assume that $f(\|p_i - q\|)$ is the cost incurred by agent i to “service” an event taking place at point q . Events take place inside Q with likelihood ϕ . Accordingly, the multicenter function $\mathcal{H}_{\text{multicenter}}$ quantifies how well the environment Q is partitioned and how well the agents are placed inside Q . This and related optimal sensor placement problems are studied in locational and geometric optimization, spatial resource allocation, quantization theory, clustering analysis, and statistical pattern recognition; see [1, Chapter 2] and references therein. \square

Among all possible ways of partitioning a subset of \mathbb{R}^d , there is one which is worth of special attention. Define the *set of partly coincident locations* $S_N = \{p \in Q^N \mid p_i = p_j \text{ for some } i, j \in \{1, \dots, N\}, i \neq j\}$. Given $p \in Q^N \setminus S_N$, the *Voronoi partition of Q generated by p* , denoted by $V(p)$, is the collection of the *Voronoi regions* $\{V_i(p)\}_{i=1}^N$, defined by

$$V_i(p) = \{q \in Q \mid \|q - p_i\| \leq \|q - p_j\| \text{ for all } j \neq i\}. \quad (2)$$

In other words, the Voronoi partition is a map $V : (Q^N \setminus S_N) \rightarrow \mathcal{V}_N$. The regions $V_i(p)$, $i \in \{1, \dots, N\}$, are convex and, if Q is a polytope, they are polytopes. Now, given two distinct points q_1 and q_2 in \mathbb{R}^d , define the $(q_1; q_2)$ -*bisector half-space* by

$$H_{\text{bs}}(q_1; q_2) = \{q \in \mathbb{R}^d \mid \|q - q_1\| \leq \|q - q_2\|\}. \quad (3)$$

In other words, $H_{\text{bs}}(q_1; q_2)$ is the closed half-space containing q_1 whose boundary is the hyperplane bisecting¹ the segment from q_1 to q_2 . Note that $H_{\text{bs}}(q_1; q_2) \neq H_{\text{bs}}(q_2; q_1)$ and that Voronoi partition of Q satisfies $V_i(p_1, \dots, p_n) = Q \cap (\cap_{j \neq i} H_{\text{bs}}(p_i; p_j))$.

Each region equipped with a density function possesses a point with a special relationship with the multi-center function. Given A , a measurable subset of Q , for each $p \in Q$ we define the scalar *1-center function* \mathcal{H}_1 by

$$\mathcal{H}_1(p; A) = \int_A f(\|p - q\|) \phi(q) dq. \quad (4)$$

One can show that, under the stated assumptions on the performance function f , the function $p \mapsto \mathcal{H}_1(p; A)$ is strictly convex in p , for any set A with positive measure (Lemma VI.1). Since this function is strictly convex, it has a unique minimum in Q . Therefore, we define the *generalized centroid* of A by

$$\text{Cd}(A) = \text{argmin}\{\mathcal{H}_1(p; A) \mid p \in Q\}. \quad (5)$$

¹A hyperplane bisects a segment if it is perpendicular to and passes through the midpoint of the segment.

In what follows, it is convenient to drop the word “generalized,” and to denote by $\text{Cd}(v) = (\text{Cd}(v_1), \dots, \text{Cd}(v_N)) \in Q^N$ the vector of regions centroids corresponding to a partition $v \in \mathcal{V}_N$.

Remark II.3 (Quadratic and linear performance functions) If the performance function is $f(x) = x^2$, then the global minimum of \mathcal{H}_1 is the *centroid* (also called the center of mass) of A , defined by

$$\text{Cd}(A) = \left(\int_A \phi(q) dq \right)^{-1} \int_A q \phi(q) dq.$$

If the performance function is $f(x) = x$, then the global minimum of \mathcal{H}_1 is the *median* (also called the Fermat–Weber center) of A . See [1, Chapter 2] for more details. \square

Proposition II.4 (Properties of $\mathcal{H}_{\text{multicenter}}$) For any partition $v \in \mathcal{V}_N$ and any point set $p \in Q^N \setminus S_N$,

$$\mathcal{H}_{\text{multicenter}}(V(p), p) \leq \mathcal{H}_{\text{multicenter}}(v, p), \quad (6)$$

$$\mathcal{H}_{\text{multicenter}}(v, \text{Cd}(v)) \leq \mathcal{H}_{\text{multicenter}}(v, p). \quad (7)$$

Furthermore, inequality (6) is strict if any entry of $V(p)$ differs from the corresponding entry of v by a set with non-empty interior, and inequality (7) is strict if $\text{Cd}(v)$ differs from p .

These statements, proved in [1, Propositions 2.14 and 2.15], motivate the following definition: a partition $v^* \in \mathcal{V}_N$ is a *centroidal Voronoi partition* if $v^* = V(\text{Cd}(v^*))$. Based on the multicenter function, we define $\mathcal{H}_{\text{centroid}} : \mathcal{V}_N \rightarrow \mathbb{R}_{\geq 0}$ by

$$\begin{aligned} \mathcal{H}_{\text{centroid}}(v) &= \mathcal{H}_{\text{multicenter}}(v, \text{Cd}(v)) \\ &= \sum_{i=1}^N \int_{v_i} f(\|q - \text{Cd}(v_i)\|) \phi(q) dq. \end{aligned} \quad (8)$$

The novel function $\mathcal{H}_{\text{centroid}}$ plays a key role in later developments and has the following property that is an immediate consequence of Proposition II.4: given a partition v with $\text{Cd}(v) \notin S_N$,

$$\mathcal{H}_{\text{centroid}}(V(\text{Cd}(v))) \leq \mathcal{H}_{\text{centroid}}(v), \quad (9)$$

and this inequality is strict if any entry of $V(\text{Cd}(v))$ differs from the corresponding entry of v by a set with non-empty interior.

B. Distributed coverage control and its limitations

Given a network of robots, coverage control algorithms move the robots in order to minimize $\mathcal{H}_{\text{multicenter}}$. To discuss these algorithms, we introduce a useful graph. The *Delaunay graph* [16], [1] associated to the distinct positions $p \in Q^N \setminus S_N$ is the undirected graph with node set $\{p_i\}_{i=1}^N$ and with the following edges: (p_i, p_j) is an edge if and only if $V_i(p) \cap V_j(p)$ is non-empty. In other words, two agents are neighbors if and only if their Voronoi regions intersect.

The *distributed coverage algorithm* presented in [1] is described as follows. At each discrete time instant $t \in \mathbb{Z}_{\geq 0}$, each agent i performs the following tasks: (i) it transmits its position and receives the positions of its neighbors in the Delaunay graph; (ii) it computes its Voronoi region with




Fig. 1. The Voronoi partition and the corresponding Delaunay graph

the information received; (iii) it moves to the centroid of its Voronoi region. In mathematical terms, for $t \in \mathbb{Z}_{\geq 0}$,

$$p(t+1) = \text{Cd}(V(p(t))). \quad (10)$$

Because of the smoothness of the various maps, compactness of Q , and monotonicity properties in Proposition II.4, one can show [1] that the solutions of (10) converge asymptotically to the set of the centroidal Voronoi partitions. This distributed coverage algorithm requires *synchronized* and *reliable* communication along *all edges of a Delaunay graph*.

This paper aims to reduce the reliability, synchronization and communication requirements of distributed coverage algorithms. Relevant questions are: Is it possible to optimize agents positions and environment partition with asynchronous, unreliable, delayed communication? What if the communication model is that of *gossiping agents*, that is, a model in which only a pair of robots can communicate at any time? How do we overcome the limitation that Voronoi partitions generated by moving agents can not be computed with only asynchronous pairwise communication?

III. PARTITIONS-BASED GOSSIP COVERAGE ALGORITHM

In the partitions-based approach, the position of the robot essentially plays no role anymore and we instead describe how to update the dominance regions. Designing coverage algorithms as dynamical systems on partitions has an important advantage: we do not restrict our attention only to Voronoi partitions.

A. The gossip coverage algorithm

Here we present a novel partition-based coverage algorithm that, at each iteration, requires only a pair of adjacent regions to communicate. We adopt the following convention: we allow communication between adjacent regions. The following definition generalizes the notion of Delaunay graph and the notion of dual graph of a planar graph [16].

Definition III.1 (Adjacency graph of a partition)

Given a partition $v \in \mathcal{V}_N$, its adjacency graph $\mathcal{G}(v)$ is the undirected graph with node set v (or equivalently $\{1, \dots, N\}$) and with the edge set $\mathcal{E}(v)$ defined as follows: (v_i, v_j) is an edge if and only if the two regions v_i and v_j are adjacent in the sense that $\text{int}(v_i) \cap \text{int}(v_j)$ is non-empty.

Recalling the notion of bisector half-space from equation (3), the *gossip coverage algorithm* is stated as follows.

At each time $t \in \mathbb{Z}_{\geq 0}$, each agent i maintains in memory a dominance region $v_i(t)$. The collection $\{v_1(0), \dots, v_N(0)\}$ is an arbitrary N -partition of Q . At each $t \in \mathbb{Z}_{\geq 0}$ a communicating pair, say $(i, j) \in \mathcal{E}(v(t))$, is selected by a deterministic or stochastic process to be determined. Every agent $k \notin \{i, j\}$ sets $v_k(t+1) = v_k(t)$, whereas agents i and j perform the following tasks:

- 1: agent i transmits to agent j its dominance region $v_i(t)$ and vice-versa
- 2: both agents compute the centroids $\text{Cd}(v_i(t))$ and $\text{Cd}(v_j(t))$
- 3: **if** $\text{Cd}(v_i(t)) = \text{Cd}(v_j(t))$ **then**
- 4: $v_i(t+1) := v_i(t)$ and $v_j(t+1) := v_j(t)$
- 5: **else**
- 6: $v_i(t+1) := (v_i(t) \cup v_j(t)) \cap H_{\text{bs}}(\text{Cd}(v_i(t)); \text{Cd}(v_j(t)))$
 $v_j(t+1) := (v_i(t) \cup v_j(t)) \cap H_{\text{bs}}(\text{Cd}(v_j(t)); \text{Cd}(v_i(t)))$
- 7: **end if**

In other words, when two agents with distinct centroids communicate, their dominance regions evolve as follows: the union of the two dominance regions is divided into two new dominance regions by the hyperplane bisecting the segment between the two centroids; see Figure 2. As a consequence, if the centroids $\text{Cd}(v_i(t))$, $\text{Cd}(v_j(t))$ are distinct, then $\{v_i(t+1), v_j(t+1)\}$ is the Voronoi partition of the set $v_i(t) \cup v_j(t)$ generated by the centroids $\text{Cd}(v_i(t))$ and $\text{Cd}(v_j(t))$. We claim that the algorithm is well-posed in




Fig. 2. The gossip coverage algorithm. The left and right figure contain the initial partition and the partition after one application of the gossip coverage algorithm. In the middle figure we show the two centroids and (with a dashed line) the segment determining the bisector half-space.

the sense that the sequence of collections $v(t)$ generated by the algorithm is an N -partition at all times t , that is, satisfies the three properties in Definition II.1.

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

$$T_{ij}(v) = \begin{cases} v, & \text{if } \text{Cd}(v_i) = \text{Cd}(v_j), \\ (v_1, \dots, \hat{v}_i, \dots, \hat{v}_j, \dots, v_N), & \text{otherwise,} \end{cases}$$

where

$$\begin{aligned} \hat{v}_i &= (v_i \cup v_j) \cap H_{\text{bs}}(\text{Cd}(v_i); \text{Cd}(v_j)), \\ \hat{v}_j &= (v_i \cup v_j) \cap H_{\text{bs}}(\text{Cd}(v_j); \text{Cd}(v_i)). \end{aligned}$$

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

$$v(t+1) = T_{ij}(v(t)), \quad \text{for some } (i, j) \in \mathcal{E}(v(t)), \quad (11)$$

together with a rule describing what edge (i, j) is selected at each time. We also define the set-valued map $T : \mathcal{V}_N \rightrightarrows \mathcal{V}_N$ by $T(v) = \{T_{ij}(v) \mid (i, j) \in \mathcal{E}(v)\}$. The map T describes

one iteration of the gossip coverage algorithm; an evolution of the gossip coverage algorithm is one of the solutions to the non-deterministic set-valued dynamical system

$$v(t+1) \in T(v(t)). \quad (12)$$

B. Designing a continuous gossip coverage algorithm

The gossip coverage map T does not satisfy certain continuity properties. In general given two regions v_i and v_j , problems arise either when $\|\text{Cd}(v_i) - \text{Cd}(v_j)\| \rightarrow 0$ or when v_i and v_j share a piece of boundary whose length tends to 0. In our convergence analysis continuity properties are necessary. Therefore, we introduce a minor modification of the gossip coverage map T , which does possess the needed continuity properties, as stated in Theorem VI.2.

Given $v = \{v_1, \dots, v_N\} \in \mathcal{V}_N$, consider two regions v_i and v_j such that $\text{Cd}(v_i) \neq \text{Cd}(v_j)$. Pick $\delta > 0$ and define²

$$\begin{aligned} \beta(v_i, v_j) &= \text{sat}_\delta(\|\text{Cd}(v_i) - \text{Cd}(v_j)\|) \\ &\quad \cdot (1 - \text{sat}_\delta(\text{dist}(\text{int}(v_i), \text{int}(v_j)))) \end{aligned}$$

where $\text{sat}_\delta : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ is such that $\text{sat}_\delta(x) = x/\delta$ if $x \in [0, \delta]$, and $\text{sat}_\delta(x) = 1$ if $x > \delta$. We aim to define a “smoothed” map T_{ij}^δ , parameterized by δ , with the following properties. If $\beta(v_i, v_j) = 1$, i.e., if the distance between the regions v_i and v_j is zero (v_i and v_j are adjacent) and the distance between $\text{Cd}(v_i)$ and $\text{Cd}(v_j)$ is larger than δ , then $T_{ij}^\delta(v) = T_{ij}(v)$; in this case the map T_{ij}^δ reduces to the map T_{ij} introduced in the previous section. Additionally, if $\beta(v_i, v_j) = 0$, i.e., either $\text{Cd}(v_i)$ and $\text{Cd}(v_j)$ coincide or the distance between v_i and v_j is larger than δ , then $T_{ij}^\delta(v) = v$, that is, the map T_{ij}^δ leaves the regions unchanged.

To define such a map T_{ij}^δ , we proceed as follows, see Figure 3. Define $R_i = v_i \cap H_{\text{bs}}(\text{Cd}(v_j), \text{Cd}(v_i)) = \{q \in v_i \mid \|q - \text{Cd}(v_i)\| \geq \|q - \text{Cd}(v_j)\|\}$, and similarly $R_j = v_j \cap H_{\text{bs}}(\text{Cd}(v_i), \text{Cd}(v_j))$. Let γ_\perp be the hyperplane bisecting the segment from $\text{Cd}(v_i)$ to $\text{Cd}(v_j)$, that is,

$$\gamma_\perp = \{q \in Q \mid \|q - \text{Cd}(v_j)\| = \|q - \text{Cd}(v_i)\|\}.$$

Observe that for each $q \in R_i \cup R_j$ there exists only one hyperplane γ parallel to γ_\perp and passing through q ; we denote this hyperplane as γ_q . Now define two points $\hat{p}_i \in R_i$ and $\hat{p}_j \in R_j$ such that

$$\begin{aligned} \hat{p}_i &\in \underset{q \in \text{int}(R_i)}{\text{argmax}} \min_{y \in \gamma_\perp} \|q - y\|, \\ \hat{p}_j &\in \underset{q \in \text{int}(R_j)}{\text{argmax}} \min_{y \in \gamma_\perp} \|q - y\|, \end{aligned}$$

and define two sets

$$\begin{aligned} \hat{R}_i &= \{q \in R_i \mid \text{dist}(\hat{p}_i, \gamma_q) \leq \beta(v_i, v_j) \text{dist}(\hat{p}_i, \gamma_\perp)\}, \\ \hat{R}_j &= \{q \in R_j \mid \text{dist}(\hat{p}_j, \gamma_q) \leq \beta(v_i, v_j) \text{dist}(\hat{p}_j, \gamma_\perp)\}. \end{aligned}$$

Now, we introduce the map $T_{ij}^\delta : \mathcal{V}_N \rightarrow \mathcal{V}_N$,

$$T_{ij}^\delta(v) = \begin{cases} v, & \text{if } \text{Cd}(v_i) = \text{Cd}(v_j), \\ (v_1, \dots, \hat{v}_i, \dots, \hat{v}_j, \dots, v_N), & \text{otherwise,} \end{cases}$$

²Given two subsets A and B of Q , define $\text{dist}(A, B) = \inf\{\|a - b\| \mid (a, b) \in A \times B\}$.




Fig. 3. Demonstration of the map T_{ij}^δ , for two regions which are close to each other.

where

$$\hat{v}_i = (v_i \setminus \text{int}(\hat{R}_i)) \cup \hat{R}_i, \quad \text{and} \quad \hat{v}_j = (v_j \setminus \text{int}(\hat{R}_j)) \cup \hat{R}_j.$$

Finally we define the *modified gossip coverage map* $T^\delta : \mathcal{V}_N \rightrightarrows \mathcal{V}_N$ by

$$T^\delta = \{T_{ij}^\delta(v) \mid (i, j) \in \{1, \dots, N\}^2, i \neq j\}. \quad (13)$$

C. Analysis and convergence results

Here we state the main convergence results of the present paper. We begin by characterizing a useful set.

Definition III.2 A partition $v \in \mathcal{V}_N$ is a mixed equal-centroidal and centroidal-Voronoi partition if, for all $(i, j) \in \mathcal{G}(v)$, either $\text{Cd}(v_i) = \text{Cd}(v_j)$ or (v_i, v_j) is a centroidal Voronoi partition of $v_i \cup v_j$.

It is easy to show that the set of mixed equal-centroidal and centroidal Voronoi partitions is equal to the set of fixed points

$$\{v \in \mathcal{V}_N \mid v = T_{ij}(v) \text{ for all } i, j \in \{1, \dots, N\}, j \neq i\}.$$

Remark III.3 Let $v \in \mathcal{V}_N$. If $\{v_i, v_j\}$ is a centroidal Voronoi partition of $v_i \cup v_j$ for any $(i, j) \in \mathcal{E}(v)$, then v is a centroidal Voronoi partition. \square

Before stating the convergence results for the modified gossip coverage algorithm, we introduce one last notion. We say that (1) the boundary of a set is *degenerate* if it has positive measure; and (2) a partition has degenerate boundary if at least one component of the partition has degenerate boundary. Note that, if each component of a partition is a collection of polygons with a finite number of vertices, then the partition boundary is not degenerate.

We now state the main deterministic convergence result for the gossip coverage algorithm.

Theorem III.4 (Convergence under uniformly persistent gossip) Given $\delta > 0$, consider the modified gossip coverage algorithm T^δ defined in (13) and let $v : \mathbb{Z}_{\geq 0} \rightarrow \mathcal{V}_N$ be an evolution of T^δ . Assume that

- (i) the trajectory v does not have an accumulation point with degenerate boundary; and
- (ii) for each pair $(i, j) \in \{1, \dots, N\}^2$, $i \neq j$, there exists an increasing sequence of times $\{t_k\}_{k \in \mathbb{Z}_{\geq 0}} \subset \mathbb{Z}_{\geq 0}$ such that $(t_{k+1} - t_k)$ is bounded and $v(t_k + 1) = T_{ij}^\delta(v(t_k))$.

Then the trajectory v converges to the set of the mixed equal-centroidal and centroidal-Voronoi partitions.

We now state the main stochastic convergence result for the gossip coverage algorithm.

Theorem III.5 (Convergence under persistent random gossip) Given $\delta > 0$, consider the modified gossip coverage algorithm T^δ defined in (13). Given a stochastic process $J : \mathbb{Z}_{\geq 0} \rightarrow \{(i, j) \in \{1, \dots, N\}^2 \mid i \neq j\}$, consider an evolution $v : \mathbb{Z}_{\geq 0} \rightarrow \mathcal{V}_N$ of T^δ satisfying, for $t \in \mathbb{Z}_{\geq 0}$,

$$v(t+1) = T_{J(t)}^\delta(v(t)).$$

Assume that

- (i) the trajectory v does not have an accumulation point with degenerate boundary; and
- (ii) there exists $p \in]0, 1[$ and $k \in \mathbb{N}$ such that, for all $(i, j) \in \{1, \dots, N\}^2$, $i \neq j$, and for all $t \in \mathbb{Z}_{\geq 0}$, there exists $h \in \{1, \dots, k\}$ such that

$$\mathbb{P}[J(t+h) = (i, j) \mid J(t), \dots, J(1)] \geq p.$$

Then the trajectory v almost surely converges to the set of the mixed equal-centroidal and centroidal-Voronoi partitions.

The proof of these two theorems is based upon the following basic result and three more complex sets of ideas. First, the basic result is a monotonicity property that clarifies the relationship between the multicenter function $\mathcal{H}_{\text{centroid}}$ and the gossip coverage algorithms.

Lemma III.6 Let $v \in \mathcal{V}_N$, $i, j \in \{1, \dots, N\}$, $i \neq j$, and $\beta \in \mathbb{R}_{>0}$. Then the gossip coverage map has the following property: $\mathcal{H}_{\text{centroid}}(T_{ij}(v)) \leq \mathcal{H}_{\text{centroid}}(v)$, and $\mathcal{H}_{\text{centroid}}(T_{ij}(v)) = \mathcal{H}_{\text{centroid}}(v)$ if and only if $T_{ij}(v) = v$. Additionally, the same result holds for the modified gossip coverage map, that is, $\mathcal{H}_{\text{centroid}}(T_{ij}^\delta(v)) \leq \mathcal{H}_{\text{centroid}}(v)$, and $\mathcal{H}_{\text{centroid}}(T_{ij}^\delta(v)) = \mathcal{H}_{\text{centroid}}(v)$ if and only if $T_{ij}^\delta(v) = v$.

This lemma indicates how the function $\mathcal{H}_{\text{centroid}}$ plays the role of a Lyapunov function for the dynamical system defined by T or T^δ . However, to provide a complete Lyapunov convergence proof, one needs to develop three sets of relevant results. First, we need to establish extensions of the LaSalle invariance principle for set-valued dynamical systems over compact metric spaces. Second, we need to prove that the space of partitions is a compact metric space. Third, we need to establish the continuity properties of the relevant maps and

functions. These three topics are the subjects of Section IV, V and VI, respectively.

D. Simulation results and implementation remarks

We have extensively simulated the partition-based gossip coverage algorithm described by (11) on a 2-dimensional polygonal environment with uniform density and performance function $f(x) = x^2$. Simulations have been implemented as a *Matlab* program, using the *General Polygon Clipping* library to perform operations on polygons. At each iteration, one edge is chosen, uniformly at random, among the edges belonging to the current adjacency graph. From these simulations, the effectiveness of the algorithm above introduced appears evident: all solutions converge to a *centroidal Voronoi partition* (Figure 4).

Our first numerical finding is that, although it is theoretically possible to converge to partitions containing regions with coincident centroids, such events do not happen in simulations. Specifically, our numerically-computed sequences of partitions always converge to centroidal Voronoi partitions, as does the synchronous coverage algorithm (10).

A second numerical finding is that, throughout numerous sample executions, regions rarely have complicated shapes and large numbers of vertices. This is good news because large numbers of vertices affect both the computation and the communication burden of the gossip coverage algorithm.

Finally, it is possible, and we have observed it numerically, to have evolutions of the algorithm that, before converging to centroidal Voronoi partitions, have components with disconnected regions. From an applications' point of view, a connected region can be covered by an agent in a more natural way. This reason suggests keeping the regions connected when applying the algorithm. We simulated the following modification of the gossip coverage algorithm which keeps the dominance regions connected: during the update step, every connected component is traded between the interacting regions only if this can be done without losing connectivity. Our simulations show that such an algorithm leads to centroidal Voronoi partitions as well.

E. A robotic implementation of gossip coverage algorithms

Consider a group of agents all having the following capabilities: (C1) each agent $i \in \{1, \dots, N\}$ knows its positions and moves at positive speed u_i to any position in the compact convex environment $Q \subset \mathbb{R}^2$; (C2) each agent may store an arbitrary number of locations in Q and has a clock that is not necessarily synchronized with other agents' clocks; and (C3) if any two agents are within distance r_{comm} of each other for a positive time duration t_{comm} and they have not communicated during the immediately prior interval of time of duration t_{comm} , then there is a positive probability p_{comm} that they establish a communication link and exchange information. It is realistic to assume $t_{\text{comm}} \ll \text{diam}(Q)/(Nu_i)$ for each i .

The *random destination+wait motion algorithm* is described as follows. Given a parameter $\epsilon < r_{\text{comm}}/4$, each agent $i \in \{1, \dots, N\}$ maintains in memory a dominance region v_i and determines its motion by repeatedly performing certain actions over periods of time that we label *epochs*. An epoch is the amount of time that agent i requires to perform the following three actions:

- 1: it selects uniformly randomly a *destination point* q_i in the set $\{q \in \mathbb{R}^2 \mid \text{dist}(q, \partial v_i \setminus \partial Q) \leq \epsilon\}$;
- 2: it moves in such a way as to reach point q_i in time equal to $t_i = \text{diam}(Q)/(N \min\{u_1, \dots, u_N\})$; and
- 3: it waits at point q_i for a time duration that is uniformly randomly distributed in the interval $[t_i, 3t_i]$.

We have assumed that agents may move outside Q and reach locations at a distance up to ϵ away from Q . This assumption may be removed at the cost of additional notation.

The random destination+wait motion algorithm is to be implemented concurrently with the modified gossip coverage algorithm with parameter $\delta < r_{\text{comm}}/4$. The two algorithms jointly determine the evolution of the agents positions and the evolution of the agents dominance regions as follows. If at any instant of time during any epoch, an agent i is within communication range r_{comm} of any other agent j for a duration t_{comm} , then, with a probability p_{comm} , the two agents exchange sufficient information to update their respective regions v_i and v_j via the modified gossip coverage map T_{ij}^δ .

Proposition III.7 (Random destination+wait ensures persistent random gossip) *Consider a group of agents with capacities (C1), (C2) and (C3) and parameters u_i , r_{comm} , t_{comm} , and p_{comm} . Assume the agents implement the random destination+wait motion algorithm and the modified gossip coverage algorithm with parameter $\epsilon < r_{\text{comm}}/4$ and $\delta < r_{\text{comm}}/4$. Then, the sequence of applications of the modified gossip coverage map satisfies the “persistent random gossip assumption” in Theorem III.5 (Assumption (ii)). Therefore, if the generated trajectory does not have an accumulation point with degenerate boundary, then the set of dominance regions maintained by the agents converges to the set of mixed equal-centroidal and centroidal-Voronoi partitions.*

IV. LASALLE INVARIANCE PRINCIPLE FOR SET-VALUED MAPS ON METRIC SPACES

In this section we consider discrete-time continuous-space set-valued dynamical system defined on metric spaces. Our goal is to provide some extensions of the classical LaSalle invariance principle; we refer the interested reader to [17], [18] for recent Lasalle invariance principles for switched continuous-time and hybrid systems.

We start by reviewing some preliminary notions. Consider a metric space (X, d) , where X is a set and d a metric on X . A set-valued map $T : X \rightrightarrows X$ is non-empty if $T(x) \neq \emptyset$ for all $x \in X$. An evolution of the dynamical system determined by a non-empty set-valued map T over X is a sequence $\{x_n \mid n \in \mathbb{Z}_{\geq 0}\} \subset X$ with the property that

$$x_{n+1} \in T(x_n), \quad n \in \mathbb{Z}_{\geq 0}.$$

Given any initial $x_0 \in X$, an evolution of T is computed by recursively setting x_{n+1} equal to an element in $T(x_n)$.

A set W is *weakly positively invariant* for T if, for any $x_0 \in W$, there exists $x \in T(x_0)$ such that $x \in W$. A set W is *strongly positively invariant* for T if, for any $x_0 \in W$, all $x \in T(x_0)$ satisfy $x \in W$.

The following result is a version of the LaSalle invariance principle for a particular class of switching dynamical systems.




Fig. 4. Simulations of randomized gossip algorithm, with $N = 6$, $Q \subset \mathbb{R}^2$, $\phi \equiv 1$ and $f(x) = x^2$. The figure shows snapshots of a time evolution of the partitions, for $t = 0, 20, 50, 100, 300$. Remark that the dominance regions can loose connectivity during the evolution.

Theorem IV.1 (Uniformly persistent switches) Let (X, d) be a 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 $\{x_n \mid n \in \mathbb{Z}_{\geq 0}\}$ be an evolution of T . Assume that:

- (i) there exists a set $W \subseteq X$ that is strongly positively invariant for T and whose closure is compact;
- (ii) there exists a function $U : W \rightarrow \mathbb{R}$ such that $U(w') < U(w)$, for all $w \in W$ and $w' \in T(w) \setminus \{w\}$;
- (iii) the functions T_i , for $i \in \{1, \dots, m\}$, and U are continuous on W ; and
- (iv) 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 $x_0 \in W$, then there exists $c \in \mathbb{R}$ such that the evolution x_n approaches the set

$$((F_1 \cap \dots \cap F_m) \cup (\partial W \setminus W)) \cap U^{-1}(c),$$

where, for $i \in \{1, \dots, m\}$, $F_i = \{w \in W \mid T_i(w) = w\}$ is the set of fixed points of the map T_i in W .

We also provide a probabilistic version of Theorem IV.1.

Theorem IV.2 (Persistent random switches) Let (X, d) be a 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 $J : \mathbb{Z}_{\geq 0} \rightarrow \{1, \dots, m\}$, consider an evolution $\{x_n \mid n \in \mathbb{Z}_{\geq 0}\}$ of T satisfying

$$x_{n+1} = T_{J(n)}(x_n).$$

Assume that:

- (i) there exists a set $W \subseteq X$ that is strongly positively invariant for T and whose closure is compact;
- (ii) there exists a function $U : W \rightarrow \mathbb{R}$ such that $U(w') < U(w)$, for all $w \in W$ and $w' \in T(w) \setminus \{w\}$;
- (iii) the functions T_i , for $i \in \{1, \dots, m\}$, and U are continuous on W ; and
- (iv) there exists $p \in]0, 1[$ 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

$$\mathbb{P}[J(n+h) = i \mid J(n), \dots, J(1)] \geq p.$$

If $x_0 \in W$, then there exists $c \in \mathbb{R}$ such that almost surely the evolution x_n approaches the set

$$((F_1 \cap \dots \cap F_m) \cup (\partial W \setminus W)) \cap U^{-1}(c),$$

where, for $i \in \{1, \dots, m\}$, $F_i = \{w \in W \mid T_i(w) = w\}$ is the set of fixed points of the map T_i in W .

V. THE SPACE OF PARTITIONS

Motivated by the results in Section IV, we study a metric structure on the set of partitions; specifically, we show how the set of partitions can be regarded as a compact metric space. In this section, and only in this section, the assumptions on Q are relaxed to give more general results: we assume that $Q \subset \mathbb{R}^d$ is compact and connected and has non-empty interior.

Let \mathcal{C} denote the set of the closed subsets of Q . Additionally, a set $C \in \mathcal{C}$ is said to be *regularly closed* if $\overline{\text{int}(C)} = C$. Given a closed set $C \in \mathcal{C}$, we say $\overline{\text{int}(C)}$ to be its *regularization*. We want to introduce a suitable metric and topology on \mathcal{C} ; since the cost functions defined in Section II are insensitive to sets of zero measure, we look for a metric with the same property.

Let μ be the Lebesgue measure of a subset of \mathbb{R}^d . Given two subsets $A, B \in \mathcal{C}$, define their *symmetric difference* by $A \Delta B = (A \cup B) \setminus (A \cap B)$ and their *symmetric distance* $d_\Delta : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ by

$$d_\Delta(A, B) = \mu(A \Delta B).$$

In other words, the symmetric distance is the measure of the symmetric difference of the two sets. Given these definitions, it is useful to identify sets that differ by a set of measure zero. More formally, let us write $A \sim B$ whenever $d_\Delta(A, B) = 0$, and remark that \sim is an equivalence relationship. In what follows we will study the *quotient set of closed subsets* $\mathcal{C}^* = \mathcal{C} / \sim$. The next result is the main result of this section.

Theorem V.1 (Metric structure and compactness of \mathcal{C}^*)

The pair $(\mathcal{C}^*, d_\Delta)$ is a metric space. Moreover, with the topology induced by the metric d_Δ , the set \mathcal{C}^* is compact.

Next, we characterize the metric structure and compactness of the set of partitions. The space of partitions \mathcal{V}_N , introduced in Section II, is mapped by the canonical projection into a \mathcal{V}_N^* , whose components belong to \mathcal{C}^* . The metric d_Δ naturally extends to a metric on the product space $(\mathcal{C}^*)^N$ and on \mathcal{V}_N^* as follows. The symmetric distance on partitions $d_\Delta : \mathcal{V}_N \times \mathcal{V}_N \rightarrow \mathbb{R}_{\geq 0}$ is defined by

$$d_\Delta(u, v) = \sum_{i=1}^N d_\Delta(u_i, v_i). \quad (14)$$

The compactness of the space of partitions is then a simple consequence of Theorem V.1.

Corollary V.2 (Metric structure and compactness of \mathcal{V}_N^*)

The pair $(\mathcal{V}_N^*, d_\Delta)$ is a metric space. Moreover, with the

topology induced by the metric d_Δ , the closure of \mathcal{V}_N^* is a compact set.

In the rest of the paper, \mathcal{V}_N^* and \mathcal{V}_N are treated as one and the same: one may think to \mathcal{V}_N as the space of the actual dynamics for the agents, and \mathcal{V}_N^* as a space which is introduced for analysis purposes. Note that, thanks to the definition of \mathcal{V}_N , \mathcal{V}_N^* can as well be depicted as a space of “partitions” made of regularly closed sets, representing the actual regions, in the sense that they differ by a set of measure zero. In general, the equivalence classes of closed sets can not be treated by means of regularly closed representatives, since the regularization of a closed set can differ from it by a set of positive measure. However, the identification can be done for closed sets satisfying the assumptions in Definition II.1, since they have zero-measure boundary.

It can be checked that all considered functions and maps of \mathcal{C} or \mathcal{V}_N are independent of the chosen representative and depend only on the equivalence class, that is, all such functions and maps are defined up to sets of measure zero. Thus, not only a sequence in \mathcal{V}_N is mapped into a sequence in \mathcal{V}_N^* , but the dynamics in \mathcal{V}_N induces a dynamics in \mathcal{V}_N^* ; it is the latter dynamics that we are able to study. Some additional useful equivalence properties are stated as follows.

Corollary V.3 *Let $u, v \in \mathcal{V}_N$ and $d_\Delta(u, v) = 0$. Then*

- (i) *the adjacency graphs of u and of v are equal;*
- (ii) *u and v have the same regularization; and*
- (iii) *if each set in u and v is regularly closed, then $u = v$.*

VI. CONTINUITY PROPERTIES OF RELEVANT MAPS

The following lemma states some important properties of the one-center cost function.

Lemma VI.1 *Let Q be the environment, and ϕ and f be the density and performance functions, respectively. For $p \in Q$, and A a compact subset of Q with positive measure, let $\mathcal{H}_1(p, A) = \int_A f(\|p - q\|)\phi(q)dq$ as in equation (4). Then*

- (i) *the map $p \mapsto \mathcal{H}_1(p, A)$ is strictly convex in p , for any A ;*
- (ii) *the map $p \mapsto \mathcal{H}_1(p, A)$ is globally Lipschitz in p , for any A ; and*
- (iii) *the map $A \mapsto \mathcal{H}_1(p, A)$ is globally Lipschitz in A , for any p .*

This lemma is a key step to prove the following results.

Theorem VI.2 (Continuity) (i) *The centroid map $Cd : \mathcal{C}^* \setminus \{\emptyset\} \rightarrow Q$, as defined in equation (5), is continuous.*
(ii) *The Voronoi map $V : Q^N \setminus S_N \rightarrow \mathcal{V}_N$, as defined in equation (2), is continuous.*
(iii) *The function $\mathcal{H}_{\text{centroid}} : \mathcal{V}_N \rightarrow \mathbb{R}_{\geq 0}$, as defined in equation (8), is continuous.*
(iv) *For all $\delta > 0$, $(i, j) \in \{1, \dots, N\}^2$, $i \neq j$, the modified gossip coverage map $T_{ij}^\delta : \mathcal{V}_N \rightarrow \mathcal{V}_N$, as defined in Section III-B, is continuous.*

Statements (iii) and (iv) are exactly what is needed to apply the LaSalle invariance principles stated in Section IV to the modified gossip coverage algorithm. Statements (i) and (ii) are intermediate results of independent interest.

VII. CONCLUSIONS

In summary, we have introduced novel multiagent coverage and partitioning algorithms, established novel versions of the LaSalle Invariance Principle, studied the topology of the space of partitions and the continuity of certain multicenter functions. Further research will focus on gossiping agents model and partition-based approaches to coverage control. First, we are keen on extending these ideas to non-convex environments: indeed parts of our analysis hold with the weaker assumption of Q being compact. Second, discrete environments like metric graphs are interesting. Third, we plan to investigate gossip coverage algorithms for robotic networks with agent arrivals and departures.

REFERENCES

- [1] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*. Applied Mathematics Series, Princeton University Press, 2009. To appear. Available at <http://www.coordinationbook.info>.
- [2] S. Martínez, J. Cortés, and F. Bullo, “Motion coordination with distributed information,” *IEEE Control Systems Magazine*, vol. 27, no. 4, pp. 75–88, 2007.
- [3] S. P. Lloyd, “Least squares quantization in PCM,” *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982. Presented as Bell Laboratory Technical Memorandum at a 1957 Institute for Mathematical Statistics meeting.
- [4] M. Pavone, E. Frazzoli, and F. Bullo, “Distributed policies for equitable partitioning: Theory and applications,” in *IEEE Conf. on Decision and Control*, (Cancún, México), pp. 4191–4197, Dec. 2008.
- [5] R. Graham and J. Cortés, “Asymptotic optimality of multicenter Voronoi configurations for random field estimation,” *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 153–158, 2009.
- [6] A. Arsie, K. Savla, and E. Frazzoli, “Efficient routing algorithms for multiple vehicles with no explicit communications,” *IEEE Transactions on Automatic Control*, 2009. To appear.
- [7] J. MacQueen, “Some methods for the classification and analysis of multivariate observations,” in *Proceedings of the Fifth Berkeley Symposium on Mathematics, Statistics and Probability* (L. M. Le Cam and J. Neyman, eds.), vol. I, pp. 281–297, University of California Press, 1965-1966.
- [8] I. I. Hussein and D. M. Stipanovič, “Effective coverage control for mobile sensor networks with guaranteed collision avoidance,” *IEEE Transactions on Control Systems Technology*, vol. 15, no. 4, pp. 642–657, 2007.
- [9] 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.
- [10] C. H. Caicedo-Núñez and M. Zefran, “Performing coverage on nonconvex domains,” in *IEEE Conf. on Control Applications*, (San Antonio, TX), pp. 1019–1024, Sept. 2008.
- [11] 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.
- [12] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Randomized gossip algorithms,” *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, 2006.
- [13] V. de Silva and R. Ghrist, “Coordinate-free coverage in sensor networks with controlled boundaries via homology,” *International Journal of Robotics Research*, vol. 25, no. 12, pp. 1205–1222, 2006.
- [14] A. Tahbaz-Salehi and A. Jadbabaie, “Distributed coverage verification in sensor networks without location information,” *IEEE Transactions on Automatic Control*, Aug. 2008. Submitted.
- [15] R. W. Brockett, “System theory on group manifolds and coset spaces,” *SIAM Journal on Control*, vol. 10, no. 2, pp. 265–284, 1972.
- [16] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry: Algorithms and Applications*. Springer, 2 ed., 2000.
- [17] A. Bacciotti and L. Mazzi, “An invariance principle for nonlinear switched systems,” *Systems & Control Letters*, vol. 54, no. 11, pp. 1109–1119, 2005.
- [18] R. G. Sanfelice, R. Goebel, and A. R. Teel, “Invariance principles for hybrid systems with connections to detectability and asymptotic stability,” *IEEE Transactions on Automatic Control*, vol. 52, no. 12, pp. 2282–2297, 2007.