

Distributed Coverage Verification in Sensor Networks Without Location Information

Alireza Tahbaz-Salehi and Ali Jadbabaie, *Senior Member, IEEE*

Abstract—In this paper, we present three distributed algorithms for coverage verification in sensor networks with no location information. We demonstrate how, in the absence of localization devices, simplicial complexes and tools from algebraic topology can be used in providing valuable information about the properties of the cover. Our approach is based on computation of homologies of the Rips complex corresponding to the sensor network. First, we present a decentralized scheme based on Laplacian flows to compute a generator of the first homology, which represents coverage holes. Then, we formulate the problem of localizing coverage holes as an optimization problem for computing a sparse generator of the first homology. Furthermore, we show that one can detect redundancies in the sensor network by finding a sparse generator of the second homology of the cover relative to its boundary. We demonstrate how subgradient methods can be used in solving these optimization problems in a distributed manner. Finally, we provide simulations that illustrate the performance of our algorithms.

Index Terms—Combinatorial Laplacians, coverage, homology, sensor networks.

I. INTRODUCTION

RECENT advances in computing, communication, sensing, and actuation technologies have brought networks composed of hundreds or even thousands of inexpensive mobile sensing platforms closer to reality. This has induced a significant amount of interest in development of analytical tools for predicting the behavior, as well as controlling the complexities of such large-scale sensor networks. Designing algorithms for deployment, localization, duty-cycling, communication, and coverage verification in sensor networks form the core of this active area of research.

Of the most fundamental problems in this domain is the *coverage problem*. In general, this reflects how well a region of interest is monitored or tracked by sensors. In most applications, we are interested in reliable monitoring of the environment in such a way that there are no gaps in the coverage. Different

algorithms for this purpose have been extensively studied [1]. One of the most prominent approaches for addressing the coverage problem has been the *computational geometry* approach, in which standard geometric tools (such as Delaunay triangulations or Voronoi diagrams) are used to determine coverage [2]–[5]. A well-known example of utilizing this geometric approach is in solving the *Art Gallery Problem*, where one determines the number of observers necessary to cover an art gallery (or an area of interest) such that every point in the gallery is monitored by at least one observer [6], [7].

Such geometrical approaches often suffer from the drawback that they can be too expensive to compute in real-time. Moreover, in most applications, they require exact knowledge of the locations of the sensors. Although, this information can be made available in real-time by a localization algorithm or by the means of localization devices (such as GPS), it can only be used most effectively in an offline, pre-deployment analysis for large networks or when there are strong assumptions on geometrical structure of the network and the environment. This drawback becomes more evident if the network topology changes due to node mobility or sensor failure. In such cases, a continuous monitoring of the network coverage becomes prohibitive if the algorithm is too expensive to run or is sensitive to location uncertainty. Finally, localization equipments add to the cost of the network, which can be a limiting factor as the size of the network grows. Consequently, a *minimal geometry* approach for addressing these issues becomes essential.

More recently, topological spaces and their topological invariants have been used in addressing the coverage problem in the absence of geometric data, such as location or orientation [8]–[14]. One notable characteristic of these studies is the use of topological abstractions which preserve many global geometrical properties of the network while abstracting away the small scale redundant details. For instance, Ghrist and Muhammad [8] construct the Rips complex corresponding to the communication graph of the network, and use the fact that the first homology of this simplicial complex contains sufficient information about the cover. Their work is followed by [9] and [11], in which a relative homological criterion for coverage is presented. These results are further extended in [10] to networks without boundary, the pursuit-evasion problem, and *barrier coverage* in 3-D. The first steps for implementation of the above mentioned ideas as distributed algorithms are taken by Muhammad and Egerstedt [12], who show that combinatorial Laplacians are the right tools for distributed computation of homologies, and hence, can be used for decentralized coverage verification. They present a consensus-like scheme based on a dynamical system whose stability properties determine the existence of coverage

Manuscript received August 01, 2008; revised April 30, 2009, accepted October 24, 2009. First published April 08, 2010; current version published July 30, 2010. This work was supported in part by the following grants: ARO/MURI W911NF-05-1-0381, DARPA/DSO SToMP, ONR MURI N000140810747, and NSF-ECS-0347285. Recommended by Associate Editor I. Paschalidis.

A. Tahbaz-Salehi is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, 02139 USA (e-mail: alirezat@mit.edu).

A. Jadbabaie is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: jadbabai@seas.upenn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TAC.2010.2047541

holes, although it fails to locate them. This idea is further extended to time-varying networks for verification of *sweep coverage* in [14].

The contribution of this paper is twofold. First, based on the ideas in [10] and [12], we present a distributed algorithm which is capable of localizing coverage holes in a network of sensors without any metric information. More precisely, we represent coverage properties of the network by its Rips complex and show that given a generator in its first homology, the problem of finding the tightest cycle encircling a hole can be formulated as an integer programming problem. We also present conditions under which the linear programming relaxation of this combinatorial problem is exact; establishing that the problem of detecting coverage holes *without the use of any coordinate information* is efficiently solvable. This optimization-based approach is a direct generalization of network flow algorithms on graphs to simplicial complexes. Finally, we show that if subgradient methods ([15]–[17]) are used for solving the relaxed problem, the subgradient updates are distributed in nature and can be implemented in a decentralized fashion. Our approach is interdisciplinary and combines results from agreement and consensus problems in multi-agent systems ([18] and [19]), with recent advances in coverage maintenance in sensor networks using computational algebraic topology methods, and optimization techniques. This novel approach is different from algorithms presented in [20] and [21], where it is explicitly assumed that the simplicial complex is embedded on an orientable surface. It is also more general than the results in [22]: our hole detection algorithm is not limited to Rips complexes, is implementable in a distributed fashion, and does not use node coordinates.

A second contribution of the paper concerns detecting redundancies in the sensor network. Using tools from algebraic topology, we introduce a novel approach for computing a minimal set of sensors required to cover the entire domain. We formulate the problem of computing the sparsest generator of the second homology of the Rips complex relative to its boundary as an integer-programming problem and solve its LP relaxation in a distributed way, using subgradient methods. To the best of our knowledge, such an algorithm has not been proposed in any other study.

The rest of the paper is organized as follows. We present the basic setup and assumptions of our model in Section II. Section III provides a brief review on the concepts of simplicial complexes, homology, and combinatorial Laplacian operators. Section IV summarizes the results already known regarding distributed coverage verification in networks with no metric information. Section V contains our main results: we show how one can “localize” coverage holes in a location-free sensor network by solving a linear programming problem using subgradient methods. We extend this idea to a higher dimension in Section VI in order to find a sparse cover of the region. Simulations of the two algorithms are presented in Section VII. Section VIII concludes.

II. PROBLEM FORMULATION

Consider a collection of n stationary sensors, denoted by V , deployed over a region of interest $\mathcal{D} \subset \mathbb{R}^2$. The sensors are equipped with local communication and sensing capabilities:

each sensor is capable of communicating with a limited number of other sensors in its proximity, and has a bounded sensing range. Furthermore, we assume a complete absence of localization capabilities and metric information, in the sense that sensors in this network can determine neither distance nor direction. Under these assumption, we are interested in distributed algorithms for coverage verification. In particular, we are interested in verifying the existence of *coverage holes*, compute their *locations*, and detect *redundancies* in the network.

We adopt the following two frameworks as coverage models for which we present our verification algorithms:

A. Simplicial Coverage

In this framework, we assume that each sensor is capable of communicating with other sensors within a radially symmetric domain of radius r_b , called the *broadcast disk*. As for the coverage, we assume a “capture” modality in which any subset of nodes which are in pairwise communication cover their entire convex hull. In other words, the region covered by the sensors is given by

$$\mathcal{A}(V) = \bigcup \left\{ \text{conv}(Q) \mid Q \subseteq V, \max_{v_i, v_j \in Q} \|v_i - v_j\|_2 \leq r_b \right\}$$

where V is the set of sensor locations and v_i represents the the location of the i -th sensor. This model, inspired by [23], guarantees that coverage and communication capabilities of the sensors are limited and based on proximity.

B. Symmetric Coverage

Similar to the previous framework, we assume that each sensor can communicate with other agents within a distance r_b . However, instead of the capture modality, we assume that sensors cover radially symmetric areas of radius r_c , known as *coverage disks*. Thus, the region covered by the sensors is $\mathcal{U}(V) = \cup_{v \in V} U_v$, where $U_v = \{x \in \mathbb{R}^2 : \|x - v\| \leq r_c\}$ is the coverage disk corresponding to the sensor located at point v . Clearly, region of interest \mathcal{D} is completely covered if it is a subset of $\mathcal{U}(V)$. For technical reasons that will become clear in the following sections, we assume that $r_b \leq r_c\sqrt{3}$. The study of this framework is motivated by networks consisting of sensors with omni-directional communication and sensing capabilities.

In addition to the assumptions on the above mentioned frameworks, we also impose some regularity conditions on the geometry of domain \mathcal{D} : we assume that \mathcal{D} is connected and compact, with a connected and piecewise linear boundary $\partial\mathcal{D}$. Finally, to avoid boundary effects, we assume that there are sensors, known as *fence nodes*, located on $\partial\mathcal{D}$ such that each fence node is capable of communicating with its two closest neighbors on $\partial\mathcal{D}$, on either side.

In the rest of the paper, we develop the required tools and present algorithms that can verify different coverage properties for the frameworks mentioned in this section.

III. SIMPLICIAL COMPLEXES, HOMOLOGY, AND COMBINATORIAL LAPLACIANS

In this section, we provide a brief review on simplicial complexes and homologies as they are the main mathematical tools

used in this paper. A thorough treatment of the subject can be found in [24] and [25].

Given a finite set of points V , a k -simplex (or a simplex of dimension k) is an unordered set $\{v_0, v_1, \dots, v_k\} \subseteq V$ where $v_i \neq v_j$ for all $i \neq j$. A *face* of the k -simplex $\{v_0, v_1, \dots, v_k\}$ is a $(k-1)$ -simplex of the form $\{v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_k\}$ for some $0 \leq i \leq k$. Clearly, any k -simplex has exactly $k+1$ faces.

Definition 1: A *simplicial complex* X is a finite collection of simplices which is closed with respect to inclusion of faces, i.e., if $\sigma \in X$, then all faces of σ are also in X .

Roughly speaking, a simplicial complex is a generalization of a graph, in the sense that in addition to binary relations between the elements of V , it also captures higher order relations between them. For example, Fig. 1 depicts a simplicial complex consisting of 11 vertices, 14 edges, five 2-simplices, and one 3-simplex. Note that due to the requirement of closure with respect to the inclusion of the faces, a simplicial complex is different from a hypergraph, in which any subset of the power set of V is considered a hyperedge.

The dimension of a simplicial complex is the maximum dimension of any of its simplices. A *subcomplex* of X is a simplicial complex $Y \subseteq X$. A particular subcomplex of X is its k -skeleton consisting of all simplices of dimension k or less, denoted by $X^{(k)} = \{\sigma \in X : \dim \sigma \leq k\}$. Therefore, the 1-skeleton of any non-empty simplicial complex is a graph. Given a graph G , its *flag complex* $F(G)$ is the largest simplicial complex whose 1-skeleton is G ; every $(k+1)$ -clique in G defines a k -simplex in $F(G)$.

Given a simplicial complex X , two k -simplices σ_i and σ_j are *upper adjacent* (denoted by $\sigma_i \frown \sigma_j$) if both are faces of a $(k+1)$ -simplex in X . Two k -simplices are said to be *lower adjacent* (denoted by $\sigma_i \smile \sigma_j$) if both have a common face. Having defined the concept of adjacency, one can define upper and lower adjacency matrices, $A_u^{(k)}$ and $A_l^{(k)}$ respectively, in order to represent the adjacency relations between different k -simplices. The upper adjacency matrix of order zero of a simplicial complex, $A_u^{(0)}$, coincides with the well-known notion of the adjacency matrix of the graph capturing its 1-skeleton.

A. Boundary Homomorphism

Let X denote a simplicial complex. Similar to graphs, an *orientation* can be defined for X by defining an ordering on every k -simplex. We denote the k -simplex $\{v_0, \dots, v_k\}$ with an ordering by $[v_0, \dots, v_k]$. For each $k \geq 0$, define $C_k(X)$ to be the vector space whose basis is the set of oriented k -simplices of X , where a change in the orientation corresponds to a change in the sign of the coefficient as $[v_0, \dots, v_i, \dots, v_j, \dots, v_k] = -[v_0, \dots, v_j, \dots, v_i, \dots, v_k]$. We let $C_k(X) = 0$, if k is larger than the dimension of X . Therefore, by definition, elements of $C_k(X)$, called k -chains, can be written as finite formal sums $\sum_j \alpha_j \sigma_j^{(k)}$ where $\alpha_j \in \mathbb{R}$ are real coefficients and $\sigma_j^{(k)}$ are the oriented k -simplices of X .¹ Note that C_k is a finite-dimensional

¹To be more precise, this is the definition of k -chains with coefficients in \mathbb{R} . In most algebraic topology texts such as [24], k -chains are defined over integers rather than reals. In such cases, $C_k(X)$ is defined as a free abelian group with the set of oriented k -simplices as its basis. However, as in [26], we find it more convenient to define the chains over \mathbb{R} .

vector space with the number of k -simplices as its dimension. We now define the boundary map.

Definition 2: For an oriented simplicial complex X , the k -th *simplicial boundary map* is a homomorphism $\partial_k : C_k(X) \rightarrow C_{k-1}(X)$, which acts on the basis elements of its domain via

$$\partial_k[v_0, \dots, v_k] = \sum_{j=0}^k (-1)^j [v_0, \dots, v_{j-1}, v_{j+1}, \dots, v_k]. \quad (1)$$

Intuitively, the above operator maps a k -chain to its faces. For example, the boundary of a directed path in a graph (which is an oriented 1-chain) is simply the difference between its two endpoints.

Since for any finite simplicial complex $C_k(X)$ is a finite dimensional vector space for all k , ∂_k has a matrix representation. We denote the matrix representation of the k -th boundary map relative to the bases of C_k and C_{k-1} by $B_k \in \mathbb{R}^{n_{k-1} \times n_k}$, where n_k is the number of k -simplices of X . In particular, the matrix representation of the first boundary map ∂_1 is nothing but the edge-vertex incidence matrix of a graph, mapping edges (1-simplices) to vertices (0-simplices).

Finally, using (1), it is an easy exercise to show that

Lemma 1: The map $\partial_k \circ \partial_{k+1} : C_{k+1}(X) \rightarrow C_{k-1}(X)$ is uniformly zero for all $k \geq 1$.

In other words, the boundary of any k -chain has no boundary.

B. Simplicial Homology

Let X denote a simplicial complex. Consider the following two subspaces of $C_k(X)$:

$$k\text{-cycles} : \ker \partial_k = \{x \in C_k(X) : \partial_k x = 0\}$$

$$k\text{-boundaries} : \text{img} \partial_{k+1} = \{x \in C_k(X) : \exists y \text{ s.t. } x = \partial_{k+1} y\}.$$

An element in $\ker \partial_k$ is a subcomplex without a boundary and therefore represents a k -dimensional cycle. On the other hand, elements of $\text{img} \partial_{k+1}$ are boundaries of higher dimensional chains and thus, are known as k -boundaries.

The k -cycles are the basic objects that count the presence of “ k -dimensional holes” in the simplicial complex [10]. But, certainly, many of the k -cycles in X are measuring the same hole; still other cycles do not really detect a hole at all—they bound a subcomplex of dimension $k+1$ in X . We say two k -cycles ξ and η are *homologous* if their difference is a boundary: $\xi - \eta \in \text{img} \partial_{k+1}$. Therefore, as far as measuring holes is concerned, homologous cycles are equivalent [10]. Consequently, it makes sense to define the quotient vector space

$$H_k(X) = \ker \partial_k / \text{img} \partial_{k+1} \quad (2)$$

known as the k -th *homology* of X , as the proper vector space for distinguishing homologous cycles. Note that according to Lemma 1, we have $\partial_k \circ \partial_{k+1} = 0$, implying that $\text{img} \partial_{k+1}$ is a subspace of $\ker \partial_k$, and therefore, making $H_k(X)$ a well-defined vector space.²

²If we define k -chains over integers, then $\text{img} \partial_{k+1}$ becomes a normal subgroup of $\ker \partial_k$. In that case, the homology is defined as the quotient group $H_k = \ker \partial_k / \text{img} \partial_{k+1}$.




Fig. 1. A simplicial complex consisting of 11 vertices (0-simplices), 14 edges (1-simplices), 5 2-simplices, and one 3-simplex.

Roughly speaking, when constructing the homology, we are removing cycles that are boundaries of a higher order subcomplex from the set of all k -cycles, so that the remaining ones carry information about the k -dimensional holes of the complex. A more precise way of interpreting (2) is that any element of $H_k(X)$ is an equivalence class of homologous k -cycles. Moreover, it inherits the structure of a vector space in the natural way: $[\xi] + [\eta] = [\xi + \eta]$ and $c[\xi] = [c\xi]$ for $c \in \mathbb{R}$, where $[\xi]$ represents the equivalence class of all k -cycles homologous to ξ . Therefore, each non-trivial homology class³ in a certain dimension identifies a corresponding “hole” in that dimension [10]. In fact, the dimension of the k -th homology of X (known as its k -th *Betti number*) identifies the number of k -dimensional holes in X . For example, the dimension of $H_0(X)$ is the number of connected components of X , whereas the dimension of $H_1(X)$ is equal to the number of holes in its 2-skeleton.

C. Relative Homology

In some applications, one may need to compute the holes modulo some region of space, such as the boundary. The concept of relative homology is defined for this purpose.

Given a simplicial complex X and a subcomplex $A \subset X$, let $C_k(X, A)$ be the quotient vector space $C_k(X)/C_k(A)$; thus, chains in A are trivial in $C_k(X, A)$. Since the boundary map $\partial_k : C_k(X) \rightarrow C_{k-1}(X)$ takes $C_k(A)$ to $C_{k-1}(A)$, it induces a quotient boundary map $\bar{\partial}_k : C_k(X, A) \rightarrow C_{k-1}(X, A)$. It is easy to verify that subspaces defined by the kernel and image of the quotient map are well-defined and satisfy $\text{img} \bar{\partial}_{k+1} \subseteq \ker \bar{\partial}_k \subseteq C_k(X, A)$. Therefore, as before, one can define the k -th *relative homology* as the quotient vector space [24]

$$H_k(X, A) = \ker \bar{\partial}_k / \text{img} \bar{\partial}_{k+1}. \quad (3)$$

Elements of $H_k(X, A)$ are equivalence classes of homologous *relative k -cycles*. A relative k -cycle is a k -chain $\xi \in C_k(X)$ such that $\partial_k \xi \in C_{k-1}(A)$. Relative k -cycle ξ is trivial in $H_k(X, A)$ if and only if it is a *relative boundary*: $\xi = \partial_{k+1} \eta + \gamma$ for some $\eta \in C_{k+1}(X)$ and $\gamma \in C_k(A)$. Fig. 2 gives examples of trivial and non-trivial relative 1-cycles.

D. Combinatorial Laplacians

The graph Laplacian [27] has various applications in image segmentation, graph embedding, dimensionality reduction for large data sets, machine learning, and more recently, in con-

³By the trivial homology class, we mean the equivalence class of all null-homologous k -cycles on the simplicial complex.




Fig. 2. 2-skeleton of a simplicial complex X and the subcomplex $A \subset X$ consisting of all boundary vertices and edges (heavy lines). Both ξ_1 and ξ_2 (dashed lines) are relative 1-cycles, but only ξ_1 represents a non-trivial element in $H_1(X, A)$.

sensus and agreement problems in distributed control of multi-agent systems [18] and [19]. For a graph G , the Laplacian matrix is defined as $L = BB^T$ where B is the vertex-by-edge-dimensional incidence matrix of G . It is well-known that the positive semi-definite Laplacian matrix can be written in terms of the adjacency and degree matrices of G : $L = D - A$, which implies that the i -th row of the Laplacian matrix only depends on the local interactions between vertex i and its neighbors. In this subsection, we present the generalization of the Laplacian matrix to simplicial complexes and investigate its properties. The importance of these generalized Laplacian matrices (known as combinatorial Laplacians) lies in the observation that, when working with real coefficients, their null spaces span subspaces isomorphic to the homologies.

The definitions and results of this subsection can be found in [26], [28], and [29].

Definition 3: Let X be a finite oriented simplicial complex. The k -th *combinatorial Laplacian* of X is the homomorphism $\mathcal{L}_k : C_k(X) \rightarrow C_k(X)$ given by

$$\mathcal{L}_k = \partial_k^* \circ \partial_k + \partial_{k+1} \circ \partial_{k+1}^* \quad (4)$$

where ∂_k^* is the adjoint of operator ∂_k with respect to the inner product that makes the basis orthonormal.

The Laplacian operator, as defined above, is the sum of two positive semi-definite operators and therefore, any k -chain $x \in \ker \mathcal{L}_k$ satisfies

$$x \in \ker \partial_k, \quad x \perp \text{img} \partial_{k+1}.$$

In other words, the null space of the k -th combinatorial Laplacian consists of k -cycles which are orthogonal to $\text{img} \partial_{k+1}$, and therefore, are not k -boundaries. This implies that non-zero elements in $\ker \mathcal{L}_k$ are representatives of non-trivial equivalence classes of cycles in the k -th homology. This property was first observed by Eckmann [28] and is formalized in the following theorem [26].

Theorem 1: If vector spaces $C_k(X)$ are defined over \mathbb{R} , then for all k there is an isomorphism

$$H_k(X) \cong \ker \mathcal{L}_k \quad (5)$$

where $H_k(X)$ is the k -th homology of X and \mathcal{L}_k is its k -th combinatorial Laplacian. Moreover, there is an orthogonal di-

rect sum decomposition of the vector space $C_k(X)$ in the form of

$$C_k(X) = \text{img} \partial_{k+1} \oplus \ker \mathcal{L}_k \oplus \text{img} \partial_k^*$$

in which the first two summands comprise the set of k -cycles $\ker \partial_k$, and the first summand is the set of k -boundaries.

An immediate implication of the above theorem is that the dimension of null space of k -th combinatorial Laplacian operator is equal to k -th Betti number of the simplicial complex.

Note that one can use matrix representations of the boundary operators to represent the combinatorial Laplacian operators with finite dimensional matrices. We define the k -th *combinatorial Laplacian matrix* as

$$L_k = B_k^T B_k + B_{k+1} B_{k+1}^T \in \mathbb{R}^{n_k \times n_k} \quad (6)$$

where B_k is the matrix representation of ∂_k and n_k is the number of k -simplices of X . It is easy to verify that the expression for L_0 reduces to the well-known graph Laplacian matrix. Similarly, combinatorial Laplacian matrices can be represented in terms of the adjacency and degree matrices of the simplicial complex [12], [29]. More precisely, for $k > 0$

$$L_k = D_u^{(k)} - A_u^{(k)} + (k+1)I_{n_k} + A_l^{(k)} \quad (7)$$

where $A_u^{(k)}$ and $A_l^{(k)}$ are the upper and lower adjacency matrices, and $D_u^{(k)}$ represents the upper degree matrix. Equation (7) implies that the i -th row of L_k only depends on the local interactions between i -th k -simplex and its upper and lower adjacent k -simplices. This is the higher-dimensional counterpart of the locality property of the graph Laplacian.

Example 1: Consider the oriented simplicial complex depicted in Fig. 3, which consists of 6 vertices, 8 edges, and 2 triangles. Its first combinatorial Laplacian matrix is given by

$$L_1 = \begin{bmatrix} 2 & -1 & 0 & 0 & 0 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 & 0 & 0 & -1 & -1 \\ 0 & -1 & 3 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 3 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 3 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 & -1 & 2 & 0 & -1 \\ 0 & -1 & 0 & 0 & 0 & 0 & 4 & 0 \\ 0 & -1 & 1 & 0 & 0 & -1 & 0 & 3 \end{bmatrix}$$

where edges are ordered as $[v_1 v_2]$, $[v_2 v_3]$, $[v_3 v_4]$, $[v_4 v_5]$, $[v_5 v_6]$, $[v_6 v_1]$, $[v_3 v_5]$, and $[v_3 v_6]$. Null space of L_1 is a one dimensional subspace spanned by vector $[8 \ 8 \ 1 \ 1 \ 3 \ 8 \ 2 \ 5]^T$. In Fig. 3, these values are depicted as flows on the edges of the simplicial complex. Notice that dimension of $\ker L_1$ is equal to the number of 1-dimensional holes in the simplicial complex, as suggested by Theorem 1. Moreover, for any $x \in \ker L_1$ the value of the algebraic sum of the flows entering each vertex is equal to zero. This is a consequence of the fact that any element in $\ker L_1$ is also in $\ker B_1$. Finally, note that the algebraic sum of the flows over any filled-in region is also equal to zero. This is due to the fact that if x is in $\ker L_1$, then $B_2^T x = 0$, and therefore, x is orthogonal to $\text{img} B_2$.




Fig. 3. Simplicial complex of dimension 2. The edge flows correspond to the components of the eigenvector corresponding to 0 eigenvalue (i.e., the null space) of the first combinatorial Laplacian matrix.

IV. DISTRIBUTED COVERAGE VERIFICATION IN THE ABSENCE OF LOCATION INFORMATION

In this section, we present a distributed coverage verification algorithm for sensor networks with no location or distance information. Unlike computational geometry approaches to coverage, the presented algorithm does not use any metric information and instead, is based on tools from algebraic topology. The main idea is to compute the null space of the first combinatorial Laplacian of a simplicial complex corresponding to the cover, and extract the information it provides on the first homology. The contents of this section are mainly based on the works of de Silva and Ghrist [10] and Muhammad and Egerstedt [12].

A. Simplicial Coverage Framework

We first consider the simplicial coverage framework. Let $V = \{v_1, \dots, v_n\}$ denote the locations of n sensors deployed over a region $\mathcal{D} \subset \mathbb{R}^2$, satisfying the assumptions presented in Section II. Sensors are equipped with local communication capabilities, enabling them to exchange data with other agents in their proximity. More specifically, two sensors can communicate with one another if the distance between them is less than or equal to r_b . As for coverage, recall that any subset of nodes in pairwise communication cover their entire convex hull. This implies that the region covered by the sensors is given by

$$\mathcal{A}(V) = \bigcup \left\{ \text{conv}(Q) \mid Q \subseteq V, \max_{v_i, v_j \in Q} \|v_i - v_j\| \leq r_b \right\}.$$

We are interested in verifying whether all points within \mathcal{D} are monitored by the sensors, i.e., whether $\mathcal{D} \subseteq \mathcal{A}(V)$. Our assumptions of Section II regarding fence nodes guarantee that $\partial \mathcal{D} \subseteq \mathcal{A}(V)$.

Since no location information is available, we capture the coverage and communication relations of sensors combinatorially. For this purpose, we define the *Vietoris-Rips complex* corresponding to a given set of points [30].

Definition 4: Given a set of points $V = \{v_1, \dots, v_n\}$ in a finite dimensional Euclidean space and a fixed radius ϵ , the Vietoris-Rips complex of V , $\mathcal{R}_\epsilon(V)$, is the abstract simplicial complex whose k -simplices correspond to unordered $(k+1)$ -tuples of points in V which are pairwise within Euclidean distance ϵ of each other. Equivalently, the Rips complex is the flag complex of the proximity graph of V , whose edges are pairs of points $v_i, v_j \in V$ with $\|v_i - v_j\| \leq \epsilon$.

The key observation is that the Rips complex corresponding to the set of sensors contains information about the region covered by them. In particular, $\mathcal{A}(V)$ is the image of the canonical projection map $p : \mathcal{R}_{r_b}(V) \rightarrow \mathbb{R}^2$ that maps each simplex affinely onto the convex hull of its vertices in \mathbb{R}^2 ; that is, the

image of p , also known as the *Rips shadow*, is exactly the region covered by the sensors. The following theorem, proved by Chambers *et al.* [23], indicates that Rips complex is rich enough to capture the relevant topological and geometric properties of its shadow.

Theorem 2: Let V denote a finite set of points in the plane, with the corresponding Rips complex $\mathcal{R}_\epsilon(V)$. Then the induced homomorphism $p_* : \pi_1(\mathcal{R}_\epsilon(V)) \rightarrow \pi_1(\mathcal{A}(V))$ between the fundamental groups of the Rips complex and its shadow is an isomorphism.

Theorem 2 states that a cycle γ in the Rips complex is contractible if and only if its projection $p(\gamma)$ is contractible in the Rips shadow [22]. As a consequence, the first homologies of the complex and its shadow are also isomorphic, creating a necessary and sufficient condition for hole-free coverage of \mathcal{D} : every point within \mathcal{D} is monitored by some sensor, if and only if the first homology of $\mathcal{R}_{r_b}(V)$ is trivial. This observation, coupled with Theorem 1, implies that the information in null space of the first combinatorial Laplacian of $\mathcal{R}_{r_b}(V)$ is sufficient to verify coverage. In particular, $\ker L_1(\mathcal{R}_{r_b}) = \{0\}$ if and only if all 1-cycles in the Rips shadow are null-homologous.

Given the above, the coverage verification problem is reduced to the study of the null space of L_1 , leading to the following theorem, first stated and proved in [12].

Theorem 3: Linear dynamical system

$$\dot{x}(t) = -L_1 x(t), \quad x(0) = x_0 \in \mathbb{R}^{n_1} \quad (8)$$

is globally asymptotically stable if and only if $\mathcal{D} \subseteq \mathcal{A}(V)$, where $x(t)$ is a vector of dimension n_1 (the number of 1-simplices of the simplicial complex) and L_1 is the first combinatorial Laplacian matrix of the Rips complex $\mathcal{R}_{r_b}(V)$.

Note that for any initial condition $x(0)$, the trajectory $x(t); t \geq 0$ always converges to a point in $\ker L_1$. Thus, the asymptotic stability of the system is an indicator of an underlying trivial homology. In different terms, since $x^* = \lim_{t \rightarrow \infty} x(t)$ is an element in the null space of L_1 , it is a representative of a homology class of the Rips complex. Clearly, if $x^* = 0$ for all initial conditions, then the first homology of the simplicial complex consists of only a trivial class and therefore, the simplicial complex is hole-free.

The importance of stating the coverage verification criterion in terms of the first combinatorial Laplacian is not limited to the above theorem. The very specific structure of the Laplacian guarantees that update (8) is effectively a local update rule, whereby the local state value of an edge is updated using estimates from edges that are adjacent to it.⁴ The reader may also note the connection between distributed update (8) and the distributed, continuous-time consensus algorithms, in which the graph Laplacian is used in order to reach a consensus (a point in the kernel) over a connected graph [19].

In summary, in order to verify coverage in a network of fixed sensors, it is sufficient to setup distributed linear dynamical system (8) for a random initial condition and observe the asymptotic state value as $t \rightarrow \infty$. If (8) converges to zero,

⁴Recall that Rips complex is the flag complex of the proximity graph and therefore, solely depends on connectivity information. As a result, its corresponding combinatorial Laplacian matrix can be formed locally by the agents.

then the first Betti number of the Rips complex is zero, and therefore, its shadow (which is the actual region covered by the sensors) is hole-free. Conversely, if the asymptotic value of (8) is non-zero for some initial condition, then the first homology of the Rips complex is non-trivial and therefore, Theorem 2 implies the existence of a non-trivial 1-cycle in the shadow and hence, the presence of holes in the cover. Note that our assumption regarding the existence of fence nodes located on the boundary of \mathcal{D} are crucial in avoiding boundary effects. The fence nodes guarantee that if a coverage hole exists, it is located in the interior of the domain.

B. Symmetric Coverage Framework

We now consider the symmetric coverage framework, in which each sensor is capable of covering a disk of radius r_c and communicate with other agents within distance $r_b \leq r_c\sqrt{3}$. In this case, the region covered by the sensors is union of disks of radius r_c centered at the location of the sensors: $\mathcal{U}(V) = \bigcup_{v_i \in V} \{x \in \mathbb{R}^2 : \|x - v_i\| \leq r_c\}$. Similar to the previous framework, we define a combinatorial object, known as the Čech or Nerve complex, that captures the topological properties of $\mathcal{U}(V)$.

Definition 5: Given a finite collection of disks $\{U_v : v \in V\}$ with radius ϵ centered at points v_i , the Čech complex of the collection denoted by $\mathcal{C}_\epsilon(V)$ is the abstract simplicial complex whose k -simplices correspond to non-empty intersections of $k+1$ distinct elements of $\{U_v : v \in V\}$.

In other words, this complex is simply formed by associating a vertex to each disk, and then adding edges and other higher order simplices based on the overlap of the disks. The following theorem, known as the Čech Theorem or the Nerve Lemma, indicates why Čech complex captures the topological properties of the region covered by sensor footprints [31], [32].

Theorem 4 (The Čech Theorem): Given a finite collection of disks U_v centered at points $v \in V$, the Čech complex $\mathcal{C}_\epsilon(V)$ has the homotopy type of the union of the disks in the collection, $\bigcup_{v \in V} U_v$.⁵

Theorem 4 implies that Čech complex contains information about homological properties of the union of the sets. In particular, both objects have isomorphic homologies in all dimensions.⁶ Therefore, in order to verify coverage in the symmetric framework, one only needs to look at the homologies of the underlying Čech complex. If this simplicial complex has no holes, neither does the sensor cover. However, computation of this simplicial complex and hence, its homologies is not an easy task, as it requires localization of each sensor as well as distance measurements in order to verify overlaps of footprints. Furthermore, as shown in [8], the Čech complex is very fragile with respect to uncertainties in distance and location measurements.

In the absence of location information, an alternative would be to use the Rips complex instead, which can be formed

⁵The statement of this theorem holds for any collection of contractible sets when all nonempty intersections of all subcollections are contractible.

⁶Note that homotopy equivalence of two topological spaces is much stronger than having isomorphic fundamental groups, as was the case in Theorem 2. In other words, the Čech complex contains much more information about the union of disks than the Rips complex does about its shadow.

uniquely from the communication graph of the network. Despite the fact that the Rips complex is not rich enough to contain all the topological and geometric information of the Čech complex, in certain cases, it carries the relevant information about homological properties of the cover. Namely, de Silva and Ghrist [10] show that a Rips complex with parameter ϵ , \mathcal{R}_ϵ , is a subcomplex of Čech complex corresponding to disks of radius $\epsilon/\sqrt{3}$ centered at its vertices. As a result, our assumption of $r_b \leq r_c\sqrt{3}$ leads to

$$\mathcal{R}_{r_b}(V) \subseteq \mathcal{C}_{r_c}(V)$$

implying $\mathcal{A}(V) \subseteq \mathcal{U}(V)$, where $\mathcal{A}(V)$ is the shadow of the Rips complex with parameter r_b and $\mathcal{U}(V)$ is the actual region covered by the sensors in the symmetric coverage framework with coverage radius r_c . Hence, if the Rips complex with parameter r_b is hole-free, so is the sensor cover. This observation serves as a sufficient homological criterion for coverage verification.

In summary, in order to verify successful coverage in a distributed fashion, sensors need to compute the first homology of the Rips complex \mathcal{R}_{r_b} using the local neighborhood information available to them. The triviality of the first homology of this simplicial complex provides a sufficient condition for hole-free coverage of \mathcal{D} . Therefore, one can set up linear dynamical system (8) corresponding to the Rips complex with parameter r_b and observe its asymptotic behavior. Similar to the simplicial coverage framework, the asymptotic stability of this dynamical system guarantees a hole-free coverage.

As a last remark, note that (8) is an edge-dimensional dynamical system, where each element of vector $x(t)$ corresponds to a 1-simplex. However, in both frameworks, edges and all other higher order simplices are simply combinatorial objects, and the only real physical entities with computational capabilities are sensors themselves. Therefore, in order to implement (8) in a sensor network, one needs a protocol for assigning the computation required by each edge to its adjacent nodes. One such algorithm is suggested by Muhammad and Jadbabaie [13], who obtain a local representation of the Rips complex and implement the dynamical system in Theorem 3 at the node level. They also show that node-level implementation of (8) is achievable by using at the most 2-hops of communications between neighboring vertices.

V. HOLE LOCALIZATION: DISTRIBUTED COMPUTATION OF A SPARSE GENERATOR

In the previous section, we presented a coverage verification algorithm which was based on the close topological relationship between the actual cover and the Rips complex as its combinatorial representation. Although capable of verifying coverage in a distributed manner, the algorithm in Section IV does not reveal any further information (such as the location, number, or size of coverage holes) about the region covered by the sensors.

In this section, we present a distributed algorithm which is capable of “localizing” coverage holes in a sensor network with no location or metric information. By hole localization, we mean detecting tight cycles over the proximity graph of the network that encircle regions not monitored by the sensors. Similar to the previous algorithm, the results of this section are based on

the algebraic topological invariants, namely the homology, of the cover and the Rips complex of the network.⁷

The main idea behind our algorithm is to compute sparse generators of a non-trivial class of homologous 1-cycles in the first homology of the Rips complex. Such sparse generators correspond to tight cycles around coverage holes, and hence provide information about their sizes and locations. Starting from a given representative cycle of a homology class, our algorithm removes components that correspond to boundaries of 2-chains in the complex, and hence, “tightens” the representative cycle around the holes. Any non-zero point in $\ker L_1$ can potentially serve as an initial 1-cycle to be tightened around the holes. The immediate advantage of using $x \in \ker L_1$ is that one can easily compute such a point in a distributed manner as the limit of linear dynamical system (8). The following example clarifies the idea behind our algorithm.

Example 2: Consider the 2-dimensional simplicial complex depicted in Fig. 3. As was shown in Example 1, the null space of its first combinatorial Laplacian is one-dimensional. Therefore, distributed linear dynamical system (8) converges to a non-zero vector in the span of $[8 \ 8 \ 1 \ 1 \ 3 \ 8 \ 2 \ 5]^T$ for almost all initial conditions. Notice that all edges, including edges $[v_3v_4]$, $[v_3v_5]$, $[v_4v_5]$, and $[v_5v_6]$ that are not adjacent to the hole, have non-zero values asymptotically. In other words, no element of $\ker L_1$ is “tight” around the hole of the simplicial complex. Another key observation is that any $x \in \ker L_1$ can be written as a linear combination of three fundamental cycles in the 1-skeleton of the simplicial complex:

$$x = 8\alpha c_1 + 3\alpha c_2 + \alpha c_3,$$

where

$$\begin{aligned} c_1 &= [1 \ 1 \ 0 \ 0 \ 0 \ 1 \ 0 \ 1]^T \\ c_2 &= [0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 1 \ -1]^T \\ c_3 &= [0 \ 0 \ 1 \ 1 \ 0 \ 0 \ -1 \ 0]^T \end{aligned}$$

and α is some real number. Among these cycles, only the first one corresponds to the hole, whereas the other two are simply contractible cycles corresponding to boundaries of 2-simplices. Therefore, in order to find a tight cycle around the hole, one needs to subtract the right amount of null-homologous 1-cycles encircling 2-simplices (in this case, 3α and α , respectively) from x . What remains is simply a 1-cycle with non-zero values only over the edges that are adjacent to the hole. Note that this cycle is also the sparsest generator of the non-trivial element of the first homology.

Due to the simple structure of the simplicial complex in the above example, computing the tightest cycle around its hole is simple. However, once the simplicial complex becomes large, it is not immediately clear how to compute the right amount of null-homologous cycles to subtract from any given element in $\ker L_1$, and obtain a sparse representative. Moreover, in the absence of a centralized scheme, it is reasonable to assume that elements of $x \in \ker L_1$ are only known locally to the nodes. This is indeed the case if x is computed in a distributed fashion using

⁷Note that in the simplicial framework, the Rips shadow coincides with the actual cover, whereas in the symmetric framework it is only a subset of the region covered by the sensors.




Fig. 4. Simplicial complexes corresponding to the above sensor configurations are combinatorially equivalent. In both cases, the shortest cycle encircling the hole has length 4.




Fig. 5. Sparsest generator of the homology class of the cycles that encircle each hole once clockwise is not necessarily tight around both.

dynamical system (8). Therefore, we need an algorithm which is capable of finding sparse non-trivial generators of homology classes of a simplicial complex by using only local information.

Before presenting our algorithm, a few remarks are in order. It is important to keep in mind that our approach is based on simplicial complexes which are combinatorial objects. As a result, in the absence of metric information, it is not possible to identify sensors which are closer in distance to the uncovered regions. For instance, consider two different sensor configurations and the regions covered by them as depicted in Fig. 4. Even though the regions covered by the sensors are different, the two configurations are combinatorially equivalent as far as Rips complex is concerned. Therefore, in both cases, any hole localization algorithm leads to the same result.

Another scenario worth mentioning is when the corresponding simplicial complex contains multiple holes next to one another. For example, consider the situation of Fig. 5, in which two holes are “close” relative to their “sizes”.⁸ In this case, the sparsest generator of the homology class of cycles that encircle each hole once is not tight: it encircles both simultaneously, rather than each one separately. Nevertheless, the sparsest 1-cycle provides valuable information on the location and size of the hole.

A. Computing the Sparsest Generator: IP Formulation

Consider a simplicial complex X with first combinatorial Laplacian L_1 . By construction, any element in the null space of L_1 is a 1-cycle orthogonal to the subspace spanned by the boundaries of 2-simplices. In other words, $x \in \ker L_1 \subset \mathbb{R}^{n_1}$ implies $x \in \ker B_1$ and $x \perp \text{img} B_2$. Therefore, as stated in Section III, any non-zero x in the null space of the first combinatorial Laplacian is a representative element of a non-trivial

⁸By terms such as close or large, we simply mean combinatorially close (in terms of hop count) and combinatorially large (in terms of the length of the shortest cycle).

homology class of X . However, as in Example 2, x is not necessarily the sparsest representative of the homology class it belongs to. In general, given a generator x of a homology class, the sparsest generator of that class can be computed as the solution to the following integer programming optimization problem:

$$\begin{aligned} & \underset{y,z}{\text{minimize}} && \|y\|_0 \\ & \text{subject to} && y = x + B_2 z \end{aligned} \quad (9)$$

where $\|\cdot\|_0$ is the ℓ_0 -norm of a vector, equal to the number of its non-zero elements, and B_2 is the matrix representation of the second boundary operator ∂_2 . Note that if x is a 1-cycle, then the minimizer y^* is also a 1-cycle in the null space of B_1 . Moreover, the constraint $y - x \in \text{img} B_2$ guarantees that both x and y^* are representatives of the same homology class, simply because adding and subtracting null-homologous cycles does not change the homology class. Therefore, any solution of the above optimization problem is the sparsest generator of the homology class that x belongs to, and has the desired property that it is the tightest possible cycle (in terms of length) around the holes represented by that homology class.

B. LP Relaxation

Optimization problem (9) has a very simple formulation. However, due to the 0–1 combinatorial element in the problem statement, solving it is not, in general, computationally tractable. In fact, Chen and Freedman [33] show that computing the sparsest generator of an arbitrary homology class is NP-hard.

A popular relaxation for solving such a problem is to minimize the ℓ_1 -norm of the objective function rather than its ℓ_0 -norm [34]

$$\begin{aligned} & \underset{y,z}{\text{minimize}} && \|y\|_1 \\ & \text{subject to} && y = x + B_2 z. \end{aligned} \quad (10)$$

This relaxation is a linear programming (LP) problem and can be solved quite efficiently.⁹ An argument similar to before shows that the minimizer of the above optimization problem is also a 1-cycle homologous to the initial x , since their difference is simply a null-homologous cycle in the image of B_2 .

In general, the minimizer of (10) is simply an approximation to the minimizer of (9) and has a larger ℓ_0 -norm. Nevertheless, in certain cases the solutions of the two problems coincide. In the next theorem, we present conditions under which the two minimizers have the same zero/non-zero pattern. Under such conditions, we would be able to compute the sparsest generator of the homology class of x efficiently.

Before formally presenting the theorem, we need to define some notation and present a lemma. Consider an oriented Rips complex \mathcal{R} with first Betti number b , where the holes are labeled 1 through b . By $h(\alpha_1, \dots, \alpha_b)$ we denote the class of homologous 1-cycles that encircle the i -th hole α_i many times in a given direction. Note that for any $\alpha \in \mathbb{R}^b$, $h(\alpha)$ is an affine subset of \mathbb{R}^{n_1} , where n_1 is the number of 1-simplices in \mathcal{R} .

⁹Strictly speaking, (10) is not a relaxation of (9), as the two problems have identical feasible sets. However, one can show, [35], that there exists an LP equivalent to (10) which is a relaxation of an IP equivalent to (9).

We assume that the shortest representative cycle encircling the i -th hole is unique, and is denoted by c_i^* ; that is

$$c_i^* = \arg \min_{c \in h(e_i)} \|c\|_0$$

where e_i is the i -th coordinate vector. Since c_i^* is the sparsest 1-cycle that encircles the i -th hole once, we have the following lemma:

Lemma 2: $c_i^* \in \{0, 1, -1\}^{n_1}$; that is, the elements of c_i^* belong to the set $\{0, 1, -1\}$.

We now present the main theorem of this section, characterizing sufficient conditions for the exactness of the relaxation problem.

Theorem 5: Given a Rips complex \mathcal{R} , suppose that the shortest representative cycle that encircles the i -th hole, denoted by c_i^* , is unique for all i . Also assume that for any simple loop $c \in h(\mu)$

$$\|c\|_0 \geq \sum_{i=1}^b |\mu_i| \|c_i^*\|_0 \quad \forall \mu \in \mathbb{Z}^b \quad (11)$$

with the inequality binding only if $c = \sum_{i=1}^b \mu_i c_i^*$. Then, for all $\alpha \in \mathbb{R}^b$, we have, $\arg \min_{c \in h(\alpha)} \|c\|_0 = \arg \min_{c \in h(\alpha)} \|c\|_1$.

Proof: Assume that edges in \mathcal{R} are labeled 1 through n_1 . We first prove that the two minimizers have identical zero/non-zero patterns. Given a class $h(\alpha)$, suppose that the ℓ_1 -minimizer, denoted by $y^{(1)}$, does not have the same pattern as the ℓ_0 -minimizer, $y^{(0)}$. This means that there exists an edge j_1 in the simplicial complex, over which $y^{(1)}$ has a non-zero value, whereas the ℓ_0 -minimizer does not. That is

$$y_{j_1}^{(1)} \neq 0, \quad y_{j_1}^{(0)} = 0.$$

Since $y^{(1)}$ is a 1-cycle, there exists another edge j_2 lower-adjacent to j_1 with a non-zero value. Reapplying the same argument implies that j_1 belongs to a set E of edges, all with non-zero values on $y^{(1)}$, forming a simple loop over the simplicial complex. Note that without loss of generality we can assume that the edge directions are defined such that the elements of $y^{(1)}$ are non-negative. Therefore

$$y_i^{(1)} > 0 \quad \forall i \in E.$$

Given set E , we define 1-chain $\tilde{c} \in \mathbb{R}^{n_1}$ as $\tilde{c}_i = \mathbf{1}_{\{i \in E\}}$, where $\mathbf{1}$ denotes the indicator function. Note that \tilde{c} is a 1-cycle which belongs to some homology class $h(\mu)$. We set $\gamma > 0$ to be the smallest value that edges in E take in the ℓ_1 -minimizer $y^{(1)}$; i.e.

$$\gamma = \min_{i \in E} y_i^{(1)}.$$

Finally, we define 1-cycle $y' = y^{(1)} - \gamma \tilde{c} + \gamma(\sum_{i=1}^b \mu_i c_i^*)$, for which we have

$$\begin{aligned} \|y'\|_1 &\leq \|y^{(1)} - \gamma \tilde{c}\|_1 + \gamma \sum_{i=1}^b |\mu_i| \|c_i^*\|_1 \\ &= \|y^{(1)}\|_1 - \gamma \|\tilde{c}\|_1 + \gamma \sum_{i=1}^b |\mu_i| \|c_i^*\|_1 \end{aligned}$$

$$\begin{aligned} &= \|y^{(1)}\|_1 - \gamma \|\tilde{c}\|_0 + \gamma \sum_{i=1}^b |\mu_i| \|c_i^*\|_0 \\ &< \|y^{(1)}\|_1. \end{aligned}$$

The first equality is due to the fact that we defined γ to be the smallest value that $y^{(1)}$ takes on E . In the second equality, we used that fact that \tilde{c} and all c_i^* are 1-cycles with values in $\{0, 1, -1\}$, implying that their ℓ_1 and ℓ_0 -norms are equal. Finally, the last inequality is a consequence of assumption (11).

In summary, there exists a 1-cycle y' homologous to $y^{(1)}$ with a smaller ℓ_1 -norm, which contradicts the fact that $y^{(1)}$ is the ℓ_1 -minimizer. Therefore, the two 1-cycles $\arg \min_{c \in h(\alpha)} \|c\|_0$ and $\arg \min_{c \in h(\alpha)} \|c\|_1$ have the same zero/non-zero pattern for all α . Finally, the fact that both 1-cycles belong to the same homology class, $h(\alpha)$, implies that the two must be equal. ■

The above theorem states that, under the given conditions, the ℓ_1 -minimizer is the sparsest generator of its homology class, and therefore, its non-zero entries indicate the edges of the 1-cycle that is tight around the holes. As a consequence, one can efficiently compute the set of edges adjacent to the holes, using methods known for solving LPs.

It is important to notice that Theorem 5 requires the uniqueness of the sparsest generator of each homology class in order to guarantee that the minimizers of the two problems coincide. When (11) holds but the ℓ_0 -minimizer is not unique, not only every ℓ_0 -minimizer is a solution to (10), but so is any convex combination of them. This is due to the fact that if two vectors have the same ℓ_1 -norm, then any vector in their convex hull cannot have a larger ℓ_1 -norm. In such cases, solving (10) can result in a 1-cycle in the convex hull of the minimizers of (9).

The intuition behind assumption (11) is also worth exploring. It requires the shortest representative cycle of any homology class to be equal to a linear combination of the shortest cycles encircling the holes separately. This condition is trivially satisfied when the simplicial complex has only one hole, or when the holes are far from each other relative to their sizes. Nevertheless, even when the condition is not satisfied, the solution of (10) is a relatively sparse (although not necessarily the sparsest) 1-cycle, and therefore, can be used as a good approximation for hole localization.

C. Decentralized Computation: The Subgradient Method

As mentioned before, unlike the original IP problem (9), one can convert (10) to a linear programming problem and solve it efficiently using methods known for solving LPs. However, applying the subgradient method enables us to compute the ℓ_1 -minimizer in a distributed manner [36], [37].

Consider optimization problem (10) reformulated as

$$\underset{z \in \mathbb{R}^{n_2}}{\text{minimize}} \quad \|x + B_2 z\|_1 \quad (12)$$

where n_2 is the number of 2-simplices of the simplicial complex. A subgradient for the objective function of (12) is the sign function. Therefore, the subgradient update for solving the above problem is given by

$$z^{(k+1)} = z^{(k)} - \alpha_k B_2^T \text{sgn}(B_2 z^{(k)} + x) \quad (13)$$

with initial condition $z^{(0)} = 0$. Note that z is a face-dimensional vector and the iteration updates an evaluation on the 2-simplices of the simplicial complex. The most important characteristic of (13) is that, due to the local structure of B_2 , it can be implemented in a distributed manner, if the initial x is known locally. By picking small enough step sizes α_k , it is guaranteed that the update (13) gets arbitrarily close to the optimal value [36], which under the conditions of Theorem 5 is the sparsest generator (or a convex combination of the sparsest generators if the minimizers are not unique) of the homology class of the initial 1-cycle x . In Section VII we provide simulations of this algorithm.

VI. DISTRIBUTED DETECTION OF REDUNDANT SENSORS

In the previous sections we presented a homological criterion for coverage. Namely, based on the results of [9], we argued that a sufficient condition for successful coverage is to have no holes in the flag complex of the proximity graph, i.e., the Rips complex of the network. This condition is translated into algebraic topological terms as $H_1(\mathcal{R}_{r_b}) = 0$, which means that every 1-cycle in the communication graph can be realized as the boundary of a surface built from 2-simplices of \mathcal{R}_{r_b} . We also showed that the first combinatorial Laplacian can be used to verify our homological criterion for coverage in a distributed manner.

In this section, we present a distributed algorithm which is capable of computing a sparse cover of domain \mathcal{D} and detect redundancies in the sensor network, in the absence of location information. In other words, the algorithm enables us to “turn off” redundant sensors without impinging upon the coverage integrity. As before, we formulate the problem of finding a sparse cover as an optimization problem to compute the sparsest generator of a certain homology class, and use subgradient methods to solve it in a distributed way. However, in contrast to the previous sections, we use the second homology of the Rips complex relative to its boundary. The advantage of the second relative homology lies in the fact that it is more robust to extensions and yields stronger information about the actual cover [10].

Consider Rips complex \mathcal{R} corresponding to the network of sensors deployed over region \mathcal{D} . We denote the subcomplex that is canonically identified with the fence nodes over $\partial\mathcal{D}$ with $\mathcal{F} \subset \mathcal{R}$. If the 1-cycles defined over \mathcal{F} are null-homologous—that is, if $[\mathcal{F}] = 0$ in $H_1(\mathcal{R})$ —then, the coverage is hole-free. In such a case, there exists a 2-chain which bounds \mathcal{F}

$$\forall \text{1-cycle } \beta \in C_1(\mathcal{F}), \quad \exists \alpha \in C_2(\mathcal{R}) \text{ s.t. } \beta = \partial_2 \alpha.$$

Therefore, when translated into the language of algebraic topology, such a 2-chain α , which is not necessarily unique, represents a relative 2-dimensional homology class, a certain generator in $H_2(\mathcal{R}, \mathcal{F})$. As a result, the condition for a hole-free successful coverage can be rewritten in terms of second relative homology classes:

Theorem 6: For a set of nodes V in a domain $\mathcal{D} \subset \mathbb{R}^2$ satisfying the assumptions of Section II, the sensor cover contains \mathcal{D} if there exists $[\alpha] \in H_2(\mathcal{R}, \mathcal{F})$ such that $\partial_2 \alpha \neq 0$.

This theorem is first stated and proved by de Silva and Ghrist [10]. Intuitively, 2-chain α has the appearance of “filling in”




Fig. 6. If the first homology of \mathcal{R} is non-trivial, then the second relative homology $H_2(\mathcal{R}, \mathcal{F})$ has no generator with values on the boundary. Conversely, if the second homology relative to the boundary has a non-trivial generator with a non-vanishing boundary, then $H_1(\mathcal{R}) = 0$.




Fig. 7. Eight faces of the octahedron form a non-trivial 2-cycle α such that $[\alpha] \in H_2(\mathcal{R})$. However, α has a vanishing boundary $\partial_2 \alpha = 0$, and does not correspond to a true relative 2-cycle.

\mathcal{D} with triangles composed of projected 2-simplices from \mathcal{R} . Note that relative homology $H_2(\mathcal{R}, \mathcal{F})$ captures the second homology of the quotient space \mathcal{R}/\mathcal{F} , in which all simplices in \mathcal{F} are identified. This can be done by adding a “super node” to the complex, as depicted in Fig. 6. If the Rips complex is hole-free, then the topology of this quotient space is that of a sphere, and therefore, the second relative homology $H_2(\mathcal{R}, \mathcal{F})$ has a non-trivial generator. On the other hand, if the 1-cycles defined over subcomplex \mathcal{F} are not boundaries of any 2-chain, then the relative homology has no generator with non-zero values on the boundary.

Note that the dimension of the second relative homology $H_2(\mathcal{R}, \mathcal{F})$ may be greater than one. This can happen if there exists a 2-cycle which is a generator of $H_2(\mathcal{R})$ as well as $H_2(\mathcal{R}, \mathcal{F})$, as depicted in Fig. 7. Such 2-cycles do not represent a true relative class, as they may still exist even if the fence cycle \mathcal{F} is not the boundary of any 2-chain. Hence, Theorem 6 requires the existence of a relative 2-cycle α with a non-zero boundary.

Given the above, it is easy to see that the minimal cover is simply the sparsest generator of the second homology class of \mathcal{R} relative to \mathcal{F} . Therefore, one can formulate the problem of finding the sparsest cover over \mathcal{D} as an optimization problem, simply by extending the results of the previous section to a higher dimension. The only difference lies in the fact that instead of the Rips complex corresponding to the network, we use the quotient complex $\bar{\mathcal{R}} = \mathcal{R}/\mathcal{F}$ which is obtained by identifying all the simplices of \mathcal{F} with a super node. Once this quotient simplicial complex is formed,¹⁰ we compute its second combinatorial Laplacian in a distributed manner, and by running the decentralized linear dynamical system $\dot{x}(t) = -L_2 x(t)$ with a

¹⁰Note that this object can be formed in a distributed fashion. The only requirement is that fence nodes take the local neighborhood relations of each other into account and update their values together.

random initial condition, obtain a point $x \in \ker L_2$ asymptotically. The limit of this dynamical update is a relative 2-cycle which does not vanish on the boundary, for almost all initial conditions. Once such a 2-cycle x is computed, the minimizer of the optimization problem

$$\begin{aligned} & \underset{y,z}{\text{minimize}} && \|y\|_0 \\ & \text{subject to} && y = x + B_3z \end{aligned} \quad (14)$$

represents the sparsest generator of the relative homology class that x belongs to. In the above problem, B_3 is the triangle-by-tetrahedron incidence matrix of the quotient complex \mathcal{R}/\mathcal{F} , x and y are 2-cycles and z is a 3-chain. Similar to problem (9), the constraint $y - x \in \text{img} B_3$ guarantees that y and x are homologous 2-cycles. Since (14) is NP-hard, one can instead solve its ℓ_1 -relaxation

$$\underset{z \in \mathbb{R}^{n_3}}{\text{minimize}} \quad \|x + B_3z\|_1 \quad (15)$$

which can be solved by the means of the distributed subgradient update

$$z^{(k+1)} = z^{(k)} - \alpha_k B_3^T \text{sgn} \left(B_3 z^{(k)} + x \right) \quad (16)$$

in a distributed manner.

Distributed iteration (16) leads to a sparse generator of the second relative homology, in which most 2-simplices have a corresponding value equal to zero. Any vertex that only belongs to 2-simplices with zero valuations in the optimal solution can be removed from the network, without generating a coverage hole. The next section contains simulations that demonstrate the performance of our algorithm.

VII. SIMULATIONS

In this section, we present the simulation results for the algorithms presented in Sections V and VI, for hole localization and computation of the minimal cover, respectively.

A. Hole Localization

We demonstrate the performance of our distributed hole localization algorithm with a randomly generated numerical example. Fig. 8(a) depicts the Rips shadow of a simplicial complex on $n = 81$ vertices distributed over \mathbb{R}^2 . The 2-skeleton of this simplicial complex consists of 81 vertices, 372 edges, and 66 triangles (2-simplices). As Fig. 8(a) suggests, the null space of the first combinatorial Laplacian of this Rips complex is 2-dimensional. The two non-trivial homology classes correspond to two eigenvectors of the Laplacian matrix corresponding to eigenvalue zero. We generated a point $x \in \ker L_1$ by running distributed linear dynamical system (8) with a random initial condition $x(0)$. The value of $x \in \ker L_1$ is depicted in Fig. 8(b), where the thickness of an edge is directly proportional to the magnitude of its corresponding component in x . It can be seen that all components of the generated 1-cycle in null space of L_1 are more or less of the same order of magnitude. In order to localize the two holes, we ran subgradient update (13) with a diminishing square summable but not summable step size. The edge evaluation of the 1-cycles after 1000 and 4000 iterations




Fig. 8. Subgradient methods can be used to localize holes in a distributed fashion. (a) The Rips shadow (b) Initial point in $\ker L_1$ (c) 1000 iterations (d) 4000 iterations.

are depicted in Figs. 8(c) and 8(d), respectively. These figures illustrate that after enough iterations, the subgradient method converges to a 1-cycle that has non-zero values only over cycles that are tight around the holes. In Fig. 8(d), the value of the 12 edges adjacent to the holes are 3 orders of magnitude higher than the rest.

Recall that our algorithm is only capable of finding the tightest minimal-length cycles surrounding the holes, which do not necessarily coincide with the cycles that are closer in distance to the holes. This is due to the fact that we are not using any metric information for our computations and the combinatorial relations between vertices is the only information available. Moreover, in the case that there are two minimal-length cycles surrounding the same hole (as in the upper hole in Fig. 8), any convex combination of those is also a minimizer of the LP relaxation problem (10). In such cases, the subgradient method in general converges to a point in the convex hull of the two solutions, rather than a corner solution. Note that the holes in the Rips complex are far relative to their sizes and therefore, Theorem 5 guarantees that the solution obtained by the ℓ_1 -minimization lies in the convex hull of the ℓ_0 -minimizers.

B. Computing a Sparse Cover

Fig. 9 illustrates the performance of the algorithm presented in Section VI. The Rips complex used for this simulation consists of 62 vertices, 22 of which function as fence nodes [Fig. 9(a)]. The second relative homology of this simplicial complex contains one non-trivial class of relative 2-cycles.

In order to compute a non-trivial representative of the second relative homology, we introduced an extra node, connected to all the fence nodes. We computed the second combinatorial Laplacian of the resulting complex and used the linear update $\dot{x}(t) = -L_2x(t)$ to obtain a point in the null space of L_2 . Subgradient update (16) is used to solve optimization problem




Fig. 9. Finding the minimal generator of the second relative homology $H_2(\mathcal{R}, \mathcal{F})$ leads to a minimal cover. 32 of the sensors can be turned off without generating any coverage holes. (a) The Rips shadow of the original cover (b) The Rips shadow of the minimal cover.

(14). The minimizer 2-cycle is depicted in Fig. 9(b). We have removed vertices that do not belong to any 2-simplex with a non-zero value at the optimal point. As illustrated in Fig. 9(b), 32 sensors can be removed from the network, without impinging upon the coverage integrity.

As a last remark, note that we can get an even sparser cover by removing either vertex a or b in Fig. 9(b). This is due to the fact that the generator computed using the subgradient update is a convex combination of two distinct solutions to the original integer programming problem (14). As in the earlier example, whenever the original problem has more than one minimizer, any relative 2-cycle in their convex hull is also a minimizer of its LP relaxation problem (15).

VIII. CONCLUSION

In this paper, we presented distributed algorithms for coverage verification in a sensor network, when no metric information is available. We used simplicial complexes and combinatorial Laplacians to capture topological properties of the network, and showed how simplicial homologies of the Rips complex can be used in verifying coverage. In particular, we illustrated the relationship between the kernel of the first combinatorial Laplacian of the Rips complex and the number of coverage holes. We formulated the problem of localizing coverage holes (in the sense of finding tight cycles encircling them) as an optimization problem that can be solved in a distributed fashion using subgradient methods. Along the same lines, we showed how one can compute a sparse cover, and detect redundancies in the network. We presented a subgradient update that is capable of computing a sparse generator of the second homology classes of the Rips complex relative to its boundary, in a decentralized manner, and used the minimizer to detect redundant sensors. Finally, we provided simulations to demonstrate the performance of our algorithms.

ACKNOWLEDGMENT

The authors would like to thank R. Ghrist for many helpful comments and discussions.

REFERENCES

- [1] D. Gage, "Command control for many-robot systems," *Unmanned Syst. Mag.*, vol. 10, no. 4, pp. 28–34, 1992.
- [2] H. Zhang and J. Hou, "Maintaining coverage and connectivity in large sensor networks," *Ad Hoc Wireless Peer-to-Peer Networks*, vol. 1, pp. 89–124, Mar. 2005.

- [3] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. Srivastava, "Coverage problems in wireless ad-hoc sensor network," in *Proc. IEEE INFOCOM*, 2001, pp. 1380–1387.
- [4] X.-Y. Li, P.-J. Wan, and O. Frieder, "Coverage in wireless ad-hoc sensor networks," *IEEE Trans. Computers*, vol. 52, no. 6, pp. 753–763, 2003.
- [5] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," in *Proc. IEEE Int. Conf. Robot. Autom.*, Washington, DC, 2002, vol. 2, pp. 1327–1332.
- [6] T. C. Shermer, "Recent results in art galleries," *Proc. IEEE*, vol. 80, no. 9, pp. 1384–1399, Sep. 1992.
- [7] A. Ganguli, J. Cortes, and F. Bullo, "Distributed deployment of asynchronous guards in art galleries," in *Proc. Amer. Control Conf.*, Minneapolis, MN, Jun. 2006, pp. 1416–1421.
- [8] R. Ghrist and A. Muhammad, "Coverage and hole-detection in sensor networks via homology," in *Proc. 4th Int. Symp. Inform. Processing Sensor Networks*, Los Angeles, CA, Apr. 2005, pp. 254–260.
- [9] V. de Silva, R. Ghrist, and A. Muhammad, "Blind swarms for coverage in 2-D," in *Proc. Robot.: Sci. Syst.*, Cambridge, MA, 2005, [CD ROM].
- [10] V. de Silva and R. Ghrist, "Coordinate-free coverage in sensor networks with controlled boundaries via homology," *Int. J. Robot. Res.*, vol. 25, no. 12, pp. 1205–1222, 2006.
- [11] V. de Silva and R. Ghrist, "Coverage in sensor networks via persistent homology," *Algebraic Geom. Topol.*, vol. 7, pp. 339–358, 2007.
- [12] A. Muhammad and M. Egerstedt, "Control using higher order Laplacians in network topologies," in *Proc. 17th Int. Symp. Math. Theory Networks Syst.*, Kyoto, Japan, 2006, pp. 1024–1038.
- [13] A. Muhammad and A. Jadbabaie, "Decentralized computation of homology groups in networks by gossip," in *Proc. Amer. Control Conf.*, New York, NY, Jul. 2007, pp. 3438–3443.
- [14] A. Muhammad and A. Jadbabaie, "Dynamic coverage verification in mobile sensor networks via switched higher order Laplacians," in *Proc. Robot.: Sci. Syst. III*, Atlanta, GA, 2007, [CD ROM].
- [15] N. Z. Shor, *Minimization Methods for Non-Differentiable Functions*. New York: Springer, 1985.
- [16] B. Polyak, *Introduction to Optimization*. Los Angeles, CA: Optimization Software, Inc, 1987.
- [17] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.
- [18] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, Jun. 2003.
- [19] R. Olfati-Saber, J. Fax, and R. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 96, no. 1, pp. 215–233, Jan. 2007.
- [20] S. Cabello, M. DeVos, J. Erickson, and B. Mohar, "Finding one tight cycle," in *Proc. 19th Annu. ACM-SIAM Symp. Discrete Algorithms*, 2008, pp. 527–531.
- [21] J. Erickson and K. Whittlesey, "Greedy optimal homotopy and homology generators," in *Proc. 16th Annu. ACM-SIAM Symp. Discrete Algorithms*, Vancouver, BC, Canada, 2005, pp. 1038–1046.
- [22] E. W. Chambers, J. Erickson, and P. Worah, "Testing contractibility in planar Rips complexes," in *Proc. 24th Annu. Symp. Computat. Geom.*, College Park, MD, 2008, pp. 251–259.
- [23] E. W. Chambers, V. de Silva, J. Erickson, and R. Ghrist, "Rips complexes of planar point sets," in *Proc. Discrete Computat. Geom.*, 2007 [Online]. Available: <http://www.math.uiuc.edu/~ghrist/preprints/rips-planar.pdf>
- [24] A. Hatcher, *Algebraic Topology*. Cambridge, U.K.: Cambridge Univ. Press, 2002.
- [25] J. R. Munkres, *Elements of Algebraic Topology*. Reading, MA: Addison Wesley, 1993.
- [26] A. M. Duval and V. Reiner, "Shifted simplicial complexes are Laplacian integral," *Trans. Amer. Math. Soc.*, vol. 354, no. 11, pp. 4313–4344, 2002.
- [27] R. Merris, "Laplacian matrices of graphs: A survey," *Linear Algebra Appl.*, vol. 197, pp. 143–176, 1994.
- [28] B. Eckmann, "Harmonische funktionen und randwertaufgaben einem komplex," *Commentarii Math. Helvetici*, vol. 17, pp. 240–245, 1945.
- [29] T. E. Goldberg, "Combinatorial Laplacians of Simplicial Complexes," B.S. thesis, Bard College, New York, NY, 2002.
- [30] L. Vietoris, "Über den höheren zusammenhang kompakter räume und eine klasse von zusammenhangstreuen abbildungen," *Mathematische Annalen*, vol. 97, no. 1, pp. 454–472, 1927.
- [31] R. Bott and L. Tu, *Differential Forms in Algebraic Topology*. Berlin, Germany: Springer-Verlag, 1982.

- [32] A. Björner, "Topological methods," in *Handbook of Combinatorics*, R. Graham, M. Grottschel, and L. Lovasz, Eds. Amsterdam, The Netherlands: North-Holland, 1995, vol. 2, pp. 1819–1872.
- [33] C. Chen and D. Freedman, Quantifying Homology Classes II: Localization and Stability 2007 [Online]. Available: arXiv:0709.2512v2
- [34] D. L. Donoho and J. Tanner, "Sparse nonnegative solution of underdetermined linear equations by linear programming," *Proc. Nat. Acad. Sci.*, vol. 102, no. 27, pp. 9446–9451, 2005.
- [35] A. Hassibi, J. How, and S. Boyd, "Low-authority controller design via convex optimization," *AIAA J. Guid., Control, Dyn.*, vol. 22, pp. 862–872, 1998.
- [36] S. Boyd, L. Xiao, and A. Mutapcic, Subgradient Methods, Lecture Notes Oct. 2003 [Online]. Available: <http://www.stanford.edu/class/ee392o/subgrad-method.pdf>
- [37] A. Nedic, Convex Optimization, Lecture Notes Apr. 2007 [Online]. Available: <https://netfiles.uiuc.edu/angelia/www/lectures.htm>



Alireza Tahbaz-Salehi received the B.Sc. degree from Sharif University of Technology, Tehran, Iran, in 2004 and the M.S.E. degree in electrical engineering, the M.A. degree in economics, and the Ph.D. degree in electrical and systems engineering from the University of Pennsylvania, Philadelphia, in 2006, 2008, and 2009, respectively.

He is currently a Postdoctoral Associate at the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge. His research interests are in distributed control of multi-agent systems, networked dynamical systems, social and economic networks, and game theory.



Ali Jadbabaie (SM'07) received the B.S. degree from Sharif University of Technology, Teheran, Iran, in 1995, the M.S. degree in electrical and computer engineering from the University of New Mexico, Albuquerque, in 1997, and the Ph.D. degree in control and dynamical systems from the California Institute of Technology, Pasadena, in 2001.

From July 2001 to July 2002, he was a Postdoctoral Associate with the Department of Electrical Engineering, Yale University, New Haven. Since July 2002, he has been with the Department of Electrical

and Systems Engineering and GRASP Laboratory, University of Pennsylvania, Philadelphia, where he is now the Skirkanich Associate Professor of Innovation. His research is broadly in control theory and network science, specifically, analysis, design and optimization of networked dynamical systems with applications to sensor networks, multi-vehicle control, social aggregation and other collective phenomena, such as synchronization and flocking.

Dr. Jadbabaie received the NSF Career Award, the ONR Young Investigator Award, the Best Student Paper Award (as advisor) of the American Control Conference 2007, the O. Hugo Schuck Best Paper Award of the American Automatic Control Council, and the George S. Axelby Outstanding Paper Award of the IEEE Control Systems Society.