ORIGINAL PAPER

Distributed computation of coverage in sensor networks by homological methods

P. Dłotko · R. Ghrist · M. Juda · M. Mrozek

Abstract   Recent work on algebraic-topological methods for verifying coverage in planar sensor networks relied exclusively on centralized computation: a limiting constraint for large networks. This paper presents a distributed algorithm for homology computation over a sensor network, for purposes of verifying coverage. The techniques involve reduction and coreduction of simplicial complexes, and are of independent interest. Verification of the ensuing algorithms is proved, and simulations detail the improved network efficiency and performance.

P. Dłotko · M. Juda · M. Mrozek
Institute of Computer Science, Jagiellonian University, Kraków, Poland
e-mail: Pawel.Dlotko@uj.edu.pl

M. Juda
e-mail: Mateusz.Juda@ii.uj.edu.pl

R. Ghrist
Departments of Mathematics and Electrical & Systems Engineering, University of Pennsylvania,
Philadelphia, PA, USA
e-mail: ghrist@math.upenn.edu

M. Mrozek (✉)
Division of Computational Mathematics, WSB-NLU, Nowy Sącz, Poland
e-mail: Marian.Mrozek@ii.uj.edu.pl

1 Introduction

Sensors—devices which return data tied to a location—are ubiquitous [12]. Although many sensors commonly used are stand-alone or *global* devices, there is an increasing push to network multiple *local* sensors, thanks to progress in miniaturization and wireless communications. The problem of collating distributed pieces of sensor data over a communications network is a substantial engineering challenge for which mathematical tools of a wide variety are relevant.

One simple-to-state problem of direct relevance is that of *coverage*, or how well a region is monitored by sensors. For concreteness, fix a domain $\mathcal{D} \subset \mathbb{R}^2$ and consider a finite collection $\mathcal{N}$ of sensors in $\mathcal{D}$. The nodes have two functions: they (1) sense a neighborhood of their locale in $\mathbb{R}^2$; and (2) they communicate with other nearby sensors. Both of these actions are assumed to be local in the sense that individual nodes cannot extract sensing data from or communicate data over all of $\mathcal{D}$. The problem of coverage, or more precisely, *blanket coverage*, is the question of whether there are holes in the sensor network—are there any regions in $\mathcal{D}$ which are not sensed? Other important coverage problems include *barrier coverage*, in which one wants to determine whether the region covered by a sensor network separates $\mathcal{D}$ or surrounds a critical region, and *sweeping coverage*, the time-dependent coverage problem familiar to users of robotic vacuum sweepers.

Coverage problems in sensor networks have received extensive attention with a literature whose tools derive from computational geometry [19,20,29,30], graphtheory [13], dynamical systems [7], and stochastic geometry [17]. Optimal coverage forms a distinct class of *art gallery* problems, with heavy representation in the computer science and complexity literature. Recently, several authors have turned to methods which are coordinate-free, inspired by problems involving non-localized ad hoc networks. In this class of coverage problems—where coordinates of sensor nodes cannot be assumed—available tools are less geometric. Among the techniques used in non-localized problems, algebraic topology has been recently seen to be both applicable and powerful [1,9,10,23,27,28]. Specifically, the tool used is *simplicial homology*, an algebraic-topological construct that determines global features (e.g., holes) in a simplicial complex (e.g., generated by nodes and communication links) utilizing local connectivity data (e.g., communications protocols). Homology theory is outlined in Appendix A; homological coverage criteria are reviewed in Sect. 3.

The problem this paper solves is as follows: in all the initial work on homological criteria for sensor coverage [9], the computations were centralized. All sensor nodes were required to upload connectivity data to a central server for the crucial homology computation. Although there are no fundamental obstructions to a decentralized computation of homology [24], specific methods for doing so are nontrivial.

Our contributions are as follows:

1. We provide a provably correct algorithm for distributed computation of homological coverage criteria.
2. The algorithm can recover specific generators, for use in building minimal or power-reducing covers.

3. Most crucially, the algorithm computes homology in an arbitrary coefficient system. This allows for computations over finite fields, which avoids the roundoff errors present in $\mathbb{R}$-coefficients.
4. Simulations seem to indicate that for a unit-disc graph network of points in the plane, all complexes are completely reducible, indicating that homology computation is of linear algorithmic complexity.

The last point leads us to the following:

Conjecture 1 For a Vietoris-Rips complex of points $\mathcal{N}$ in the plane $\mathbb{R}^2$, the S-complex $\mathcal{K}^f$ remaining after applying Algorithm 8 is always boundaryless.

If this conjecture is true, there is no need for any additional computations to ascertain the homology. This would be beneficial as regards theoretical complexity bounds, since the reduction algorithm in a non-distributed setting and bounded amount of neighbors has complexity $O(N)$, where N is the number of simplices (see [21]). However, even if in some situations the conjecture fails, in practice the size of the remaining complex is very small: one node would easily compute the homology of the remaining complex.

Our work compares most closely to the recent approach of Jadbabaie and Tahbaz-Salehi [27]. They also derive a distributed algorithm for homology computation with the goal of satisfying the homological coverage criterion, determining the existence/location of holes, and generating optimal coverage. Their methods involve passing to homology with $\mathbb{R}$ coefficients and setting up the problem as a dynamical system—in essence, they use simplicial Laplacians and their connection to homology via Hodge theory [11] to solve a heat equation over the network. This can be done in a distributed manner, using message-passing and gossip algorithms [16]. The paper [27] is very creative in that they also apply compressive-sensing perspectives to the problem of provably and quickly determining optimal generators for homology classes.

The reduction/coreduction scheme of this paper has several advantages.

1. It greatly reduces the communication complexity demanded by a dynamical systems approach as in [27]: waiting for a large network to asymptotically converge to a solution is nontrivial in time and in energy drain from communication.
2. It is applicable to arbitrary coefficient systems. The Hodge-theoretic approach of [27] cannot avoid the use of $\mathbb{R}$ coefficients and the ensuing roundoff errors, which can accumulate to obscure answers in settings where homology generators are not well-separated. Our approach works well with finite field coefficients (e.g., mod-2 arithmetic) or integer coefficients, avoiding roundoff altogether.

Moreover, the methods developed in this paper may be adapted to speed up general-purpose homology computations in the context of parallel architecture, in particular in multi-core processors and GPU's (work in progress). Finally, more flexible local coefficient systems (as in [14]) may be adaptable; if so, the techniques of this paper may be extended to give a distributed computation of the cohomology of simplicial sheaves, a problem whose relevance to networks is emerging [15,25,26].

We understand that an independent method for distributed homology computation is being investigated by Carlsson, de Silva, and Morozov, using the technique of *zigzag persistence*, as initiated in [4] and [5].

2 Sensor and network assumptions

For our applications of distributed homology computation to coverage problems, we operate under the following assumptions.

1. Sensors are modeled as a collection of nodes $\mathcal{N} \subset \mathbb{R}^2$.
2. Each sensor is assumed to have a unique identification which it broadcasts; certain neighbors detect the transmission and establish a communication link.
3. Communication links are symmetric, stable, and generate a well-defined communications graph $\mathcal{G}$ on $\mathcal{N}$.
4. Sensor coverage regions are correlated to communications: the convex hull of any subset of sensors $S \subset \mathcal{N}$ which pairwise communicate is contained in the union of coverage regions of S.
5. One fixes a cycle $\mathcal{C} \subset \mathcal{G}$ whose image in $\mathbb{R}^2$ is a simple closed curve bounding a simply connected domain $\mathcal{D} \subset \mathbb{R}^2$.

Under these assumptions, one wants to know whether $\mathcal{D}$ is contained in the coverage region of the network. These assumptions are chosen to be weak enough to be applicable in realistic systems; however, some important considerations—e.g., time-variability, node failure, false echoes, communications errors—are not modeled.

3 Homological coverage

This paper builds on a homological criterion for coverage in sensor networks described and explored in [9,10]. This section reviews that criterion. The reader for whom homology is foreign will want to make a brief excursion to Appendix A.

3.1 Simplicial complexes for networks

This paper uses as its basic data structure simplicial complexes based on a set of nodes (sensors) $\mathcal{N}$. A finite family X of nonempty finite subsets of $\mathcal{N}$ is an (abstract) *simplicial complex* if for every $\sigma \in X$ and $\tau \subset \sigma$ we have $\tau \in X$. The elements $\sigma \in X$ are *simplices* whose *dimension* equals the cardinality minus one: dim $\sigma = |\sigma| - 1$. The 0-simplices of X are *vertices* and the 1-simplices are *edges*, as in graph theory. A *face* of a simplex σ in a simplicial complex X is a simplex $\tau \subset \sigma$ with dim $\tau = \dim \sigma - 1$. For τ a face of σ one says that σ is a *coface* of τ. Every simplicial complex has a natural geometric realization—also denoted X—as a topological space obtained by gluing disjoint copies of the standard k-simplex, Δ^k, along faces, as in Fig. 1.

Simplicial complexes are natural to networks. Any communications network on nodes $\mathcal{N}$ has the structure of a simplicial complex with all simplices of dimension zero or one. In several applied contexts, higher dimensional simplicial complexes are natural data structures. Examples include the following. Assume a set of nodes $\mathcal{N}$ in Euclidean $\mathbb{R}^n$:

1. The size ϵ *Čech complex* is the simplicial complex $\mathcal{C}_\epsilon$ on $\mathcal{N}$ whose k-simplices are generated by $k + 1$ nodes about which the diameter ϵ balls have a mutually nonempty intersection.

Fig. 1 The geometric realization of a simplicial complex

2. The size ϵ *Vietoris-Rips complex* is the simplicial complex $\mathcal{R}_\epsilon$ on $\mathcal{N}$ whose k-simplices are generated by $k + 1$ nodes about which the diameter ϵ balls have *pairwise* nonempty intersections. Namely, simplices are tuples of nodes with pairwise distance less than or equal to ϵ.

3. Given a network of edges based on $\mathcal{N}$, the *flag complex* of the network is the simplicial complex $\mathcal{F}$ whose k simplices are generated by pairwise-connected $(k + 1)$-tuples of nodes in the network.

Thus, the Vietoris-Rips complex $\mathcal{R}_\epsilon$ is the flag complex of the size ϵ *unit disc graph*. This simple model of connectivity for a sensor or communications network is widely used [13,20,30] and widely disparaged [3,8,18]. Our results hold in the context of more general flag complexes associated to network communication graphs, whether they are unit disc graphs or not.

In this paper, we restrict attention to simplicial complexes. There is a broader notion of a *cell complex* having as building blocks not merely simplices, but cubes or other polyhedral cells. These can sometimes be useful in modeling the geometry underlying a sensor network. The methods described in this paper extend to these as well but are not detailed explicitly.

3.2 The homological coverage criterion

Recall our standing assumption linking coverage to communication: any triple of nodes in pairwise communication has its convex hull in $\mathbb{R}^2$ contained in the coverage region, as in Fig. 2. We emphasize that this neither assumes localization nor an unrealistic assumption about round balls, symmetry, or the like: this is a very flexible model. As stated, it is explicitly simplicial, and motivates the use of the flag complex $\mathcal{F}$ of the communications graph $\mathcal{G}$ to model the topology of the sensed region.

The second important assumption is that a cycle $\mathcal{C}$ in the network is chosen whose image in $\mathbb{R}^2$ is a simple closed curve bounding a polygonal domain $\mathcal{D}$. This cycle acts as a *fence* for the coverage problem. (Criteria for guaranteeing that a cycle in a non-localized network has simple image in the plane are simple to derive [6,9]—it suffices to have no 'shortcuts' between cycle nodes.) The following theorem gives a criterion for coverage based on (relative) homology of the flag complex $\mathcal{F}$ of the com-

Fig. 2 Coverage regions are assumed to be correlated to communication distances: the convex hull of a triple of communicating nodes is assumed covered
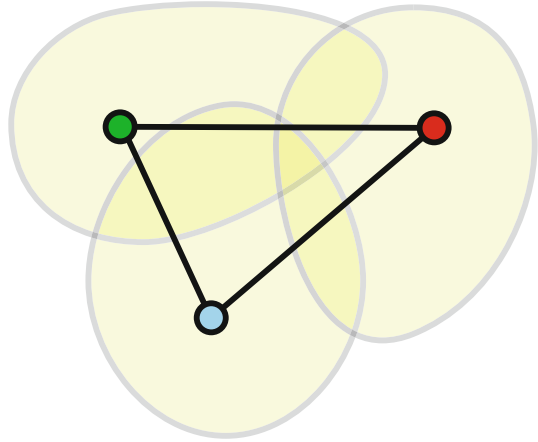
munications network modulo the fence cycle $\mathcal{C}$. It is a slight modification of Theorem 3.3 of [9].

Theorem 1 *(Homological coverage criterion* [9]*) Given a planar sensor network with $\mathcal{G}$, $\mathcal{C}$, and $\mathcal{D}$ as above, then all of $\mathcal{D}$ is completely covered by the sensors if, equivalently:*

1. $[\mathcal{C}] = 0 \in H_1(\mathcal{F})$.
2. *There exists $[\zeta] \in H_2(\mathcal{F}, \mathcal{C})$ with $\partial \zeta = \mathcal{C}$.*

In practice, the second form of the criterion is more useful, since, if an explicit relative cycle $\zeta \in Z_2(\mathcal{F}, \mathcal{C})$ is computed, then it provides a guarantee of coverage even when the nodes not implicated in ζ are removed (or 'powered down' for conservation reasons), *cf.* [9].

4 S-complexes and reduction algorithms

Our algorithm for a distributed homology computation that suffices to check the criterion of Theorem 1 necessitates a modification of the relevant chain complexes. Our approach is explicitly simplicial and uses reductions and coreductions to simplify a simplicial complex. Recalling (Appendix A) that in homology theory, a simplicial complex (a fundamentally topological object) is replaced with a chain complex (a fundamentally algebraic device), we simplify the simplicial and chain complexes simultaneously.

4.1 S-complexes

We review the concept of an *S-complex* introduced in [21] as a reformulation of a chain complex, convenient for algorithmic purposes. Let $\mathcal{K} = (K_q)_{q \in \mathbb{Z}}$ be a finite graded set, whose grading is called *dimension*, denoted dim. Let $R[\mathcal{K}]$ be the graded free module over the unitary ring R of *chains*, generated by the graded set $\mathcal{K}$, with inner product $\langle \cdot, \cdot \rangle$ induced by $\mathcal{K}$. We refer to the generators of $R[\mathcal{K}]$ as *elementary chains*.

Let $\kappa : \mathcal{K} \times \mathcal{K} \to R$ be a map for which $\kappa(\sigma, \tau) \neq 0$ only if $\dim \sigma = \dim \tau + 1$. This induces a morphism $\partial^{\kappa} : R[\mathcal{K}] \to R[\mathcal{K}]$ via

$$\partial^{\kappa}(\sigma) := \sum_{\tau \in \mathcal{K}} \kappa(\sigma, \tau)\tau \qquad \sigma \in \mathcal{K}.$$

This is well-defined and of degree -1. We say that $(\mathcal{K}, \kappa)$ is an *S-complex* if $(R[\mathcal{K}], \partial^{\kappa})$ is a chain complex. The homology of an S-complex $(\mathcal{K}, \kappa)$, denoted $H(\mathcal{K})$, is defined as the homology of the chain complex $(R[\mathcal{K}], \partial^{\kappa})$. In the sequel we will drop the superscript κ in ∂^{κ} whenever κ is clear from the context. We also refer to $R[\mathcal{K}]$ as a chain complex if ∂^{κ} is clear from the context.

Every simplicial complex gives rise to an S-complex. Indeed, every simplicial complex $\mathcal{K}$ has a natural gradation $(K_q)_{q \in \mathbb{Z}}$, where K_q consists of simplices of dimension q. Assume an ordering of the vertices of $\mathcal{K}$ is given and every simplex σ in K_q is coded as $[v_0, v_1, \ldots, v_q]$, where $v_0, v_1, \ldots, v_q$ are listed according to the prescribed ordering. By putting

$$\kappa(\sigma, \tau) := \begin{cases} (-1)^i & \text{if} \quad \sigma = [v_0, v_1, \ldots, v_{i-1}, v_i, v_{i+1}, \ldots, v_q] \\ & \text{and} \quad \tau = [v_0, v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_q] \\ 0 & \text{otherwise.} \end{cases}$$

we obtain an S-complex $(\mathcal{K}, \kappa)$.

We say that $(\mathcal{K}, \kappa)$ is *boundaryless* if $\partial^{\kappa} = 0$. Note that if $(\mathcal{K}, \kappa)$ is boundaryless, then there is no computation involved in determining homology, as $H(\mathcal{K}) = R[\mathcal{K}]$. A boundaryless complex is highly desirable from the point of view of algorithms, as no further computation is necessary. Our strategy is to incrementally and asynchronously modify S-complexes so as to achieve a boundaryless complex.

Given an elementary chain $\sigma \in \mathcal{K}$ of an S-complex, we define its *boundary* and *coboundary sets* as:

$$\mathrm{bd}_{\mathcal{K}}(\sigma) := \{ \tau \in \mathcal{K} \mid \kappa(\sigma, \tau) \neq 0 \},$$
$$\mathrm{cbd}_{\mathcal{K}}(\sigma) := \{ \sigma \in \mathcal{K} \mid \kappa(\sigma, \tau) \neq 0 \}.$$

We say that $\mathcal{K}' \subset \mathcal{K}$ is *closed* in $\mathcal{K}$ if $\mathrm{bd}_{\mathcal{K}} \mathcal{K}' \subset \mathcal{K}'$. Note that if $\mathcal{K}'$ is closed in $\mathcal{K}$, then $\partial^{\kappa}(R[\mathcal{K}']) \subset R[\mathcal{K}']$; therefore, there is a well defined restriction $\partial^{\kappa}|_{R[\mathcal{K}']} : R[\mathcal{K}'] \to R[\mathcal{K}']$, which gives rise to a chain subcomplex $(R[\mathcal{K}'], \partial^{\kappa}|_{R[\mathcal{K}']})$ of the chain complex $(R[\mathcal{K}], \partial^{\kappa})$.

We say that $\mathcal{K}' \subset \mathcal{K}$ is *open* in $\mathcal{K}$ if the complement $\mathcal{K} \setminus \mathcal{K}'$ is closed. Note that if $\mathcal{K}'$ is open in $\mathcal{K}$, then there is a well defined quotient complex $(R[\mathcal{K}]/R[\mathcal{K} \setminus \mathcal{K}'], \partial')$ with the boundary map ∂' taken as the respective quotient map of ∂^{κ}.

A subset $\mathcal{K}' \subset \mathcal{K}$ is an *S-subcomplex* of the S-complex $\mathcal{K}$ if $(\mathcal{K}', \kappa')$, with $\kappa' := \kappa|_{\mathcal{K}' \times \mathcal{K}'}$, the restriction of κ to $\mathcal{K}' \times \mathcal{K}'$, is itself an S-complex, i.e. if $(R[\mathcal{K}'], \partial^{\kappa'})$ is a chain complex.

One can easily verify that if $\mathcal{K}'$ is closed in $\mathcal{K}$, then $\partial^{\kappa'}$ coincides with $\partial^{\kappa}|_{R[\mathcal{K}']}$. In consequence, $\mathcal{K}'$ is an S-subcomplex of $\mathcal{K}$ and the homology of $\mathcal{K}'$ coincides with the homology of the chain subcomplex $(R[\mathcal{K}'], \partial^{\kappa}|_{R[\mathcal{K}']})$.

A lengthier but easy argument shows that if $\mathcal{K}'$ is open in $\mathcal{K}$, then $\partial^{\kappa'}$ is conjugate to the boundary map ∂' in the quotient complex $(R[\mathcal{K}]/R[\mathcal{K} \setminus \mathcal{K}'], \partial')$. Thus, also in this case $\mathcal{K}'$ is an S-subcomplex of $\mathcal{K}$ and the homology of $\mathcal{K}'$ is isomorphic to the homology of the quotient complex $(R[\mathcal{K}]/R[\mathcal{K} \setminus \mathcal{K}'], \partial')$.

Therefore, we have the following theorem (see also [21, Theorem 3.2]).

Theorem 2 *If $\mathcal{K}'$ is closed or open in $\mathcal{K}$, then $\mathcal{K}'$ is an S-subcomplex of $\mathcal{K}$. Moreover, if $\mathcal{K}'$ is closed in $\mathcal{K}$, then $H(\mathcal{K}')$ coincides with the homology of the chain subcomplex $(R[\mathcal{K}'], \partial^{\kappa}|_{R[\mathcal{K}']})$ and if $\mathcal{K}'$ is open in $\mathcal{K}$, then $H(\mathcal{K}')$ is isomorphic to the homology of the quotient complex $(R[\mathcal{K}]/R[\mathcal{K} \setminus \mathcal{K}'], \partial')$.*

However, the concept of an S-subcomplex is broader than the concepts of closed and open subsets, because, as the following example shows, there are *S*-subcomplexes which are neither open nor closed.

Example 1 Take a simplicial complex $\mathcal{K}$ which consists of four vertices A, B, C, D, four edges AB, AC, BC, CD and one triangle ABC. Consider the following subsets:

$$\mathcal{K}' := \mathcal{K} \setminus \{D\},$$
$$\mathcal{K}'' := \mathcal{K}' \setminus \{CD, C\},$$
$$\mathcal{K}''' := \mathcal{K}'' \setminus \{ABC, AB\}.$$

One easily verifies that $\mathcal{K}'$ is open in $\mathcal{K}$, $\mathcal{K}''$ is open in $\mathcal{K}'$ and $\mathcal{K}'''$ is closed in $\mathcal{K}''$. Therefore, $R[\mathcal{K}']$ is a quotient complex of $R[\mathcal{K}]$, $R[\mathcal{K}'']$ is a quotient complex of $R[\mathcal{K}']$ and $R[\mathcal{K}''']$ is a chain subcomplex of $R[\mathcal{K}'']$. However, $R[\mathcal{K}''']$ is neither a subcomplex nor a quotient complex of $R[\mathcal{K}']$ or $R[\mathcal{K}]$. But $\mathcal{K}'''$ is an S-subcomplex of $\mathcal{K}''$, $\mathcal{K}'$ and $\mathcal{K}$.

The following proposition is straightforward.

Proposition 1 *Assume $\mathcal{K}'$ is an S-subcomplex of an S-complex $\mathcal{K}$ and $\mathcal{A} \subset \mathcal{K}'$. Then*

$$\mathrm{bd}_{\mathcal{K}'}(\mathcal{A}) = \mathrm{bd}_{\mathcal{K}}(\mathcal{A}) \cap \mathcal{K}',$$
$$\mathrm{cbd}_{\mathcal{K}'}(\mathcal{A}) = \mathrm{cbd}_{\mathcal{K}}(\mathcal{A}) \cap \mathcal{K}'.$$

The following is an easy consequence of Theorem 2 and standard homological algebra (see Appendix A).

Theorem 3 [21, Theorem 3.4] *For $\mathcal{K}'$ closed and $\mathcal{K}''$ the complementary open subset of $\mathcal{K}$, the short exact sequence*

$$0 \to R[\mathcal{K}'] \xrightarrow{\iota} R[\mathcal{K}] \xrightarrow{\pi} R[\mathcal{K}''] \to 0$$

with inclusion ι and projection π induces the following long exact sequence of homology modules:

$$\cdots \xrightarrow{\delta} H_q(\mathcal{K}') \xrightarrow{H(\iota)} H_q(\mathcal{K}) \xrightarrow{H(\pi)} H_q(\mathcal{K}'') \xrightarrow{\delta} H_{q-1}(\mathcal{K}') \cdots \qquad (1)$$

4.2 Reduction, coreduction, and sequences

The parallelized simplification of an S-complex $\mathcal{K}$ consists of removing certain pairs of elements of $\mathcal{K}$ while leaving the homology of the complex intact. Given a pair $\alpha = (\tau, \sigma) \in \mathcal{K}^2$, we refer to the doubleton $\{\tau, \sigma\}$ as the *support* of the pair α and denote it $|\alpha|$. We extend the concept and notation of support to collections $\mathcal{B} = \{\beta_1, \beta_2, \ldots, \beta_n\} \subset \mathcal{K}^2$ of pairs by unions: $|\mathcal{B}| := \cup_{i=1}^{n} |\beta_i|$.

A pair $\alpha = (\tau, \sigma) \in \mathcal{K}^2$ is an *elementary reduction pair in* $\mathcal{K}$ if $\mathrm{cbd}_{\mathcal{K}}\{\tau\} = \{\sigma\}$ and $\kappa(\sigma, \tau)$ is invertible in R. It is called an *elementary coreduction pair in* $\mathcal{K}$ if $\mathrm{bd}_{\mathcal{K}}\{\sigma\} = \{\tau\}$ and again $\kappa(\sigma, \tau)$ is invertible in R. For a fixed S-complex $\mathcal{K}$, an *S-reduction pair* is a pair which is either an elementary reduction pair or an elementary coreduction pair in $\mathcal{K}$. Sample reductions of an S-complex performed by removing the support of an S-reduction pair are presented in Fig. 3.

Lemma 1 *Given α an S-reduction pair in an S-complex $\mathcal{K}$, we have the following properties*

(i) *If α is an elementary reduction (resp. coreduction) pair, then $|\alpha|$ is open (resp. closed) in $\mathcal{K}$.*
(ii) *The support $|\alpha|$ is an S-subcomplex of $\mathcal{K}$ and $H(|\alpha|) = 0$.*
(iii) *There is a well defined chain map*

$$R[\mathcal{K} \setminus |\alpha|] \ni c \mapsto c - \frac{\langle \partial c, \tau \rangle}{\langle \partial \sigma, \tau \rangle} \sigma \in R[\mathcal{K}] \tag{2}$$

inducing an isomorphism

$$\gamma_\alpha : H(\mathcal{K} \setminus |\alpha|) \cong H(\mathcal{K}). \tag{3}$$

(iv) *If α is an elementary reduction pair, then γ_α coincides with the isomorphism induced by the inclusion $\iota_\alpha : R[\mathcal{K} \setminus |\alpha|] \to R[\mathcal{K}]$.*
(v) *If α is an elementary coreduction pair, then γ_α coincides with the inverse of the isomorphism induced by the projection $\pi_\alpha : R[\mathcal{K}] \to R[\mathcal{K} \setminus |\alpha|]$.*

Fig. 3 Simplification of a complex (*center*) performed via elementary coreduction (*left*) and reduction (*right*) pairs

Proof Properties (i) and (ii) are direct (and follow from [21, Theorem 4.1] and Theorem 2). Properties (iii), (iv) and (v) use the long exact sequence of the pair (and follow from [22, Corollary 2.7], [22, Theorem 2.8], [22, Theorem 2.9] and Theorem 3). ☐

Note that a pair $\alpha \in \mathcal{K}^2$ which is not necessarily an S-reduction pair in $\mathcal{K}$ may be an S-reduction pair in a S-subcomplex of $\mathcal{K}$. We call such an S-reduction pair a *potential S-reduction pair in* $\mathcal{K}$. Whenever potential S-reduction pairs exist, the reduction process may be self-feeding: removing some S-reduction pairs may give rise to new S-reduction pairs. On the other hand, if the supports of two S-reduction pairs in a given S-complex have non-empty intersection, performing one of these reductions eliminates the other reduction as a possibility. For disjoint supports, both reductions may be performed independently, in parallel. We say that a collection $\mathcal{B} = \{\beta_1, \beta_2, \ldots, \beta_n\}$ of potential S-reduction pairs in $\mathcal{K}$ is *free* in an S-subcomplex $\mathcal{L}$ of $\mathcal{K}$ if all elements of $\mathcal{B}$ are S-reduction pairs in $\mathcal{L}$ and any two distinct elements of $\mathcal{B}$ have disjoint supports. We say that $\mathcal{B}$ is *free* if it is free in some S-subcomplex of $\mathcal{K}$.

Let $\mathcal{L}$ be an S-subcomplex of $\mathcal{K}$. We say that an ordered sequence $\varphi = (\varphi_1, \varphi_2, \ldots, \varphi_n)$ of potential S-reduction pairs in $\mathcal{K}$ is a *reduction sequence* in $\mathcal{L}$ if for each $j = 1, 2, \ldots, n$ the pair φ_j is an S-reduction pair in $\mathcal{L} \setminus \bigcup_{i=1}^{j-1} |\varphi_i|$. Note that a straightforward recursive argument shows that the definition is meaningful, i.e., $\mathcal{L} \setminus \bigcup_{i=1}^{j-1} |\varphi_i|$ is an S-complex. A reduction sequence is said to be *free* if the elements of the sequence form a free collection of S-reduction pairs.

We use S-reduction pairs to selectively simplify an S-complex. A reduction sequence φ in an S-subcomplex $\mathcal{L}$ of $\mathcal{K}$ yields an isomorphism

$$I_\varphi : H(\mathcal{L} \setminus |\varphi|) \to H(\mathcal{L}) \quad : \quad I_\varphi := \gamma_{\varphi_1} \circ \gamma_{\varphi_2} \circ \cdots \circ \gamma_{\varphi_n}, \tag{4}$$

given by the composition of the isomorphisms γ_{φ_i} of (3). The following result is crucial in proving correctness of our algorithms:

Theorem 4 *Let $\mathcal{L}$ be an S-subcomplex of $\mathcal{K}$. Free collections and reduction sequences have the following properties.*

(i) *Any total ordering of a free collection $\mathcal{B}$ of S-reduction pairs in $\mathcal{L}$ forms a reduction sequence.*
(ii) *The set $\mathcal{L} \setminus |\mathcal{B}|$ is an S-subcomplex of $\mathcal{K}$ for any free collection $\mathcal{B}$ in $\mathcal{L}$.*
(iii) *Two free reduction sequences φ, φ' which differ by a permutation give rise to the same isomorphism $I_\varphi = I_{\varphi'}$.*

Proof The proof of property (i) proceeds by induction in m, the size of the free collection $\mathcal{B}$. If $m = 1$, the conclusion is obvious. For $m > 1$ the conclusion follows immediately from Proposition 1 and the induction assumption. Property (ii) follows immediately from (i) and the definition of the reduction sequence. The proof of property (iii) proceeds by induction in n, the common length of the free reduction sequences φ and φ'. If $n = 1$, the conclusion is obvious. If $n = 2$, then $\varphi = (\varphi_1, \varphi_2)$ and $\varphi' = (\varphi_2, \varphi_1)$ unless $\varphi = \varphi'$ when the conclusion is obvious. There are three cases to be considered: either both φ_1 and φ_2 are elementary reduction pairs or they are both elementary coreduction pairs or one is an elementary reduction pair and the other is

an elementary coreduction pair. In the first case we have the following commutative diagram of chain maps induced by inclusions

$$R[\mathcal{K} \setminus |\varphi_1|] \longleftarrow R[\mathcal{K} \setminus (|\varphi_1| \cup |\varphi_2|)]$$

$$\downarrow \qquad\qquad\qquad \downarrow$$

$$R[\mathcal{K}] \longleftarrow R[\mathcal{K} \setminus |\varphi_2|]$$

which induces a commutative diagram in homology. By Lemma 1(iv), we obtain $I_\varphi = I_{\varphi'}$. An analogous argument based on a diagram with projections and a diagram with inclusions combined with projections bring us to the same conclusion in the other two cases. There remains to consider the case $n > 2$. Since φ and φ' differ only by a permutation, there exists an $i \in \{1, \ldots, n\}$ such that $\varphi'_i = \varphi_n$. Let

$$\bar{\varphi} := (\varphi_1, \varphi_2, \ldots, \varphi_{n-1})$$
$$\bar{\psi} := (\varphi'_1, \varphi'_2, \ldots, \varphi'_{i-1}, \varphi'_{i+1}, \ldots \varphi'_n)$$
$$\psi := (\varphi'_1, \varphi'_2, \ldots, \varphi'_{i-1}, \varphi'_{i+1}, \ldots \varphi'_n, \varphi'_i)$$

Obviously, $\bar{\varphi}$ and $\bar{\psi}$ are free reduction sequences. By the induction assumption $I_{\bar{\varphi}} = I_{\bar{\psi}}$, therefore

$$I_\varphi = I_{\bar{\varphi}} \circ \gamma_{\varphi_n} = I_{\bar{\psi}} \circ \gamma_{\varphi'_i} = I_\psi.$$

However, by the case $n = 2$ applied $n - i$ times to ψ we get $I_\psi = I_{\varphi'}$. This proves that $I_\varphi = I_{\varphi'}$. $\qquad\qquad\square$

Theorem 4 makes a distributed reduction process possible: the reductions in a free reduction sequence may be performed independently, so in parallel, independent of order.

Example 2 Consider the set of nine sensors $\{A, B, C, D, E, F, G, H, I\}$ whose communication graph consists of eighteen edges (see Fig. 4)

$$\{AB, AE, AG, BC, BE, BF, CD, CE, CF,$$
$$DF, DI, EF, EG, EH, FH, FI, GH, HI\}.$$

After removing the fence consisting of edges $\{AB, BC, DI, HI, GH, AG\}$, we obtain an S-complex composed of two vertices, eleven 1-simplices, eleven 2-simplices and one 3-simplex. A search for S-reduction pairs reveals elementary coreduction pairs (F, DF) and (E, CE). These form a free collection and may be removed in parallel. Continuing, we obtain the following reduction sequence, in which free collections are written in separate lines.

$$(F, DF)\ (E, CE)$$
$$(FI, DFI)\ (CF, CDF)\ (BE, BCE)$$
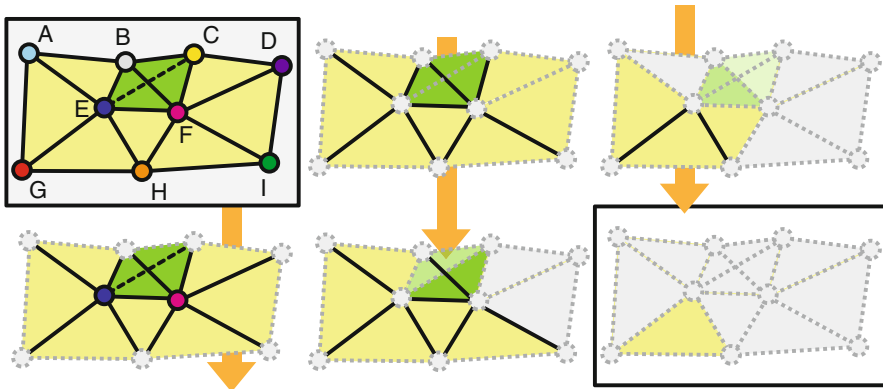
Fig. 4 From *top left*, to *bottom right*, following the *arrows*: a collection of sensors and the associated S-complex, the S-complex after removing the fence, the fence complex after removing the consecutive free collections of S-reduction pairs, the final reduced S-complex

$$(EF, CEF)\ (BF, BCF)\ (AE, ABE)\ (FH, FHI)$$
$$(BEF, BCEF)\ (EG, AEG)\ (EH, EFH).$$

In the case of this reduction sequence the resulting S-complex is boundaryless. It consists only of the (open) 2-simplex *EFH*, which becomes the generator of the second homology group of the S-complex.

5 Construction of the flag complex

In this section we describe the distributed algorithms for the construction of the flag complex as a simplicial approximation of the sensor network. Henceforth we assume a fixed enumerative ordering of the sensor nodes $\mathcal{N}$, as per unique identification numbers. We denote the total order by $\lhd$. Moreover, we assume that all algorithms described in this section are executed separately on every node.

Following the conventions of object oriented programming, we write s.alg() to indicate that node *s* is requested to execute alg(). In every algorithm we denote the node running the algorithm by the word this. We drop the prefix "this" in this.alg(), assuming that by default the node supposed to execute the called algorithm is the node making the call. We assume that every node stores the set of its neighbors in variable neighbors.

To store simplex σ of the flag complex $\mathcal{F}$ associated with sensor network $\mathcal{N}$ we use a data structure Simplex with the following fields and methods:

- σ.vert—the set of the vertices (nodes) in σ,
- σ.neighb—the set of the neighbors of σ,
- σ.faces()—returns the set of the faces of σ,
- σ.cofaces()—returns the set of the cofaces of σ.

In course of running the reduction algorithm we also need some auxiliary fields. In particular, when a simplex is removed by the reduction algorithm, instead of deleting

---

Algorithm 1 newSimplex(V, N)

1. $\sigma :=$ new simplex object.
2. $\sigma.\texttt{vert} := V$.
3. $\sigma.\texttt{neighb} := N$.
4. $\sigma.\texttt{locked} := \sigma.\texttt{deleted} := \texttt{false}$.
5. return σ.

Algorithm 2 createLocalFlagComplex

1. $\tau_0 := \texttt{newSimplex}(\{\texttt{this}\}, \texttt{neighbors})$.
2. $\texttt{Simp}[0] := \{\tau_0\}$.
3. $i := \max\{j \in \mathbb{N} \mid \texttt{Simp}[j] \neq \emptyset\}$.
4. for each $\tau \in \texttt{Simp}[i]$
   (a) next τ if $\min(\tau.\texttt{vert}) \neq \texttt{this}$.
   (b) for each $t \in \tau.\texttt{neighb}$
       i. next t if $\min(\tau.\texttt{vert}) \trianglelefteq t \trianglelefteq \max(\tau.\texttt{vert})$.
       ii. $V := \tau.\texttt{vert} \cup \{t\}$,
       iii. $N := \tau.\texttt{neighb} \cap t.\texttt{neighbors}$,
       iv. $\sigma := \texttt{newSimplex}(V, N)$.
       v. $\texttt{Simp}[i+1] := \texttt{Simp}[i+1] \cup \{\sigma\}$.
5. if $\texttt{Simp}[i+1] \neq \emptyset$, then go to step 3.

---

it from the data structure we only mark a flag indicating that the simplex is to be treated as deleted. We use another flag to lock a simplex when a node is negotiating its removal with its neighbors. Here is the list of all auxiliary fields:

- $\sigma.\texttt{deleted}$—a boolean variable marking that σ is deleted,
- $\sigma.\texttt{locked}$—a boolean variable marking that σ is locked.

In order to construct a simplex we use Algorithm 1. It returns a simplex with vertex set V and the neighbors set N.

Algorithm 2 is the algorithm for building the flag complex $\mathcal{F}$ of the network as expressed locally in terms of vertex-neighbor data. The algorithm is executed on each node. In the sequel, by $\texttt{s.Simp}$ we mean the set of simplices created by node s in the course of running Algorithm 2. The subset of $\texttt{s.Simp}$ consisting of simplices of dimension i is denoted by $\texttt{s.Simp}[i]$. We say that node $\texttt{s}$ *controls* simplex σ if simplex σ is created in course of running Algorithm 2 on node $\texttt{s}$.

Theorem 5 *Assume Algorithm 2 has been initiated on each node. Then it terminates on each node. Moreover, once the algorithm has completed on all nodes, the following properties hold true.*

1. *For every node* $\texttt{s}$ *and every simplex*

$$\sigma \in \texttt{s.Simp}$$

variables $\sigma.\texttt{vert}$ *and* $\sigma.\texttt{neighb}$ *store respectively, as expected, the vertices and neighbors of simplex* σ.

2. *Every $\sigma \in \mathcal{F}$ is controlled by at least one and at most two nodes. It is controlled by exactly two nodes if and only if* dim $\sigma > 0$. *Moreover, these two nodes are the two minimal elements of σ.* vert.

3. *The only simplices created by Algorithm 2 are simplices in $\mathcal{F}$.*

Since every simplex σ is controlled by at least one and most two nodes, in the sequel we speak about the *lower* and the *upper* node controlling σ, assuming that the lower equals the upper if only one node controls σ. We omit the definitions of σ. faces() and σ. cofaces(). They may be computed by a node using Simp stored in the lower and in the upper node controlling σ.

Proof To see that the algorithm always terminates, first observe that the loop in line (4) always completes, because S is finite. The only other loop in the algorithm is formed by the conditional jump from line (5) to line (3). Note that whenever the jump takes place, the variable i is increased. However, the value of i may not exceed the number of nodes. Therefore, the jump in line (5) may be performed only a finite number of times and consequently the algorithm terminates.

The fact that variable σ. vert stores the vertices of σ is an immediate consequence of Algorithm 1. Let $\sigma \in$ s. Simp be a simplex and let d be its dimension. We will prove by induction on d that variable σ. neighb contains the neighbors of σ. If $d = 0$, the conclusion is obvious from the construction. Thus, assume that the conclusion holds for all simplices of dimension not exceeding d. The simplex σ is created from a simplex τ and its neighbor t at line (4.b.iii) of Algorithm 2. In particular, we see from the construction that

$$\sigma.\texttt{neighb} = \tau.\texttt{neighb} \cap \{u \in V \mid \{u, t\} \in E\}.$$

However, by the induction assumption

$$\tau.\texttt{neighb} = \{u \in V \mid \forall_{v \in \tau.\texttt{vert}} \{v, u\} \in E\}$$

and the conclusion follows from elementary set arithmetic. This proves property (1).

Let us now demonstrate property (2). First consider the case when dim $\sigma = 0$. Observe that the only line of the algorithm which contains the construction of a zero dimensional simplex is line (1) of the algorithm. Moreover, it is straightforward to check that this line is executed exactly once in each node. Therefore, a zero dimensional simplex τ_0 is constructed by node s and only by node s. There remains to consider the case dim $\sigma > 0$. Let

$$\sigma.\texttt{vert} = \{v_0, v_1, \ldots, v_n\}.$$

Observe that the only line of Algorithm 2 which contains a construction of a simplex of dimension greater than zero is line (4.b.iv). Let τ denote the simplex selected in line (4) just before the execution of line (4.b.iv). Since line (4.b.iv) is executed, the test in line (4.a) fails, which implies that

$$\min(\tau.\texttt{vert}) = \texttt{this}. \tag{5}$$

Algorithm 3 canDelete(σ)

1. return this = min(σ.vert) and not σ.deleted and not
   σ.locked.

The execution of lines (4.b.ii) and (4.b.iii) implies that the difference

$$\sigma.\texttt{vert} \setminus \tau.\texttt{vert}$$

contains exactly one element t and, by the execution of line (4.b.i), we know that either
$t \lhd \min(\tau.\texttt{vert})$ or $t \rhd \max(\tau.\texttt{vert})$. The first case happens when $t = v_0, \tau = \{v_1, v_2, \ldots, v_n\}$ and the other case when $t = v_n, \tau = \{v_0, v_1, \ldots, v_{n-1}\}$. Therefore, by (5), the node running the algorithm is either v_0 or v_1. It follows that if $\dim \sigma > 0$, then σ is represented in the memory of the two minimal elements of $\sigma.\texttt{vert}$.

To prove property (3) suppose by contrary, that there exists a simplex $\sigma \in \mathcal{F}$ which has not been created by the algorithm in any node. From line (2) of the algorithm it follows that $\dim \sigma > 0$. Let σ be the simplex of minimal dimension with the described property. It follows that all the faces τ of σ have been created by the algorithm. From line (4) of the algorithm, it follows that σ must have been created by the algorithm: a contradiction.

In a similar way one can show that if $\sigma \notin \mathcal{F}$ then it cannot be created by the algorithm. Suppose by contrary, that such a σ has been created and σ is of minimal dimension among such simplices. Therefore, σ has been constructed in line (4.b.iv) of the algorithm from a simplex τ by adding a vertex $t \in \tau.\texttt{vert}$. By the minimality assumption, $\tau \in \mathcal{F}$, and since $t \in \tau.\texttt{vert}$, we get $\sigma \in \mathcal{F}$, a contradiction.   □

The following Lemma is used in the Algorithm 6.

Lemma 2 *For all simplices σ created by nodes $\texttt{s}$ and $\texttt{t}$ each face $\tau \in \sigma.\texttt{faces()}$ is created by node $\texttt{s}$ or $\texttt{t}$.*

Proof By Theorem 5 we have $\texttt{s}, \texttt{t} \in \sigma.\texttt{vert}$. For each face τ of σ, we have $\text{card}(\sigma.\texttt{vert} \setminus \tau.\texttt{vert}) = 1$, therefore $\texttt{t} \in \tau.\texttt{vert}$ or $\texttt{s} \in \tau.\texttt{vert}$. Since $\tau.\texttt{vert} \subset \sigma.\texttt{vert}$, we have $\min(\tau.\texttt{vert}) \in \{\texttt{s}, \texttt{t}\}$ and by Theorem 5 simplex τ is created by node $\texttt{s}$ or $\texttt{t}$.   □

6 Reductions

The parallel reduction procedures used in the algorithm will now be presented.

6.1 Fence reduction

We assume that Algorithm 5 is executed on each node. Definitions of the auxiliary procedures delete and canDelete are presented as Algorithm 3 and Algorithm 4.

Algorithm 4 delete(σ)

1. if $\sigma \in$ Simp and not σ.deleted, then
   (a) σ.deleted := σ.locked := true.
   (b) for each node u controlling σ
       u.delete(σ).

Algorithm 5 removeFence

1. if this $\notin F$ then exit.
2. τ_0 := the unique simplex in Simp[0].
3. delete(τ_0).
4. for each s' $\in$ neighbors and for each $\sigma \in$ s'.Simp[1] s.t.
   card $\{\tau \in \sigma$.faces() | not τ.deleted$\} = 0$
   and s'.canDelete(σ)
   (a) s'.delete(σ).
5. for each $\sigma \in$ Simp[1] s.t.
   card $\{\tau \in \sigma$.faces() | not τ.deleted$\} = 0$
   and canDelete(σ)
   (a) delete(σ).

Lemma 3 *Algorithm 5 deletes all simplices in $\mathcal{C}$ and no other simplices.*

Proof A zero-dimensional simplex τ_0 is deleted at line (3) of the algorithm and line (1) guarantees that only 0-simplices in $\mathcal{C}$ can be deleted there. A 1-simplex of $\mathcal{C}$ is deleted at line (4.a) and (5.a) of Algorithm 5. Since at this stage the only deleted 0-simplices belong to $\mathcal{C}$, it is clear that at line (4.a) and (5.a) only 1-simplices from $\mathcal{C}$ can be deleted. From the definition of the fence as a simple cycle it is clear that in the points (4.a) and (5.a) all the 1-simplices in $\mathcal{C}$ are deleted. Since lines (3), (4.a) and (5.a) are the only lines where the reduction of a simplex takes place, no other simplices can be removed. □

Since $\mathcal{C}$ is a closed subset of $\mathcal{F}$ in the sense of the definition in Sect. 4, it follows that $\mathcal{F} \setminus \mathcal{C}$, together with the usual boundary map, is an S-complex.

6.2 Distributed S-reductions

From Lemma 1, S-reduction pairs may be removed from an S-complex without changing the homology groups of the complex. We present now how the reductions may be performed in a distributed manner. We perform the reduction process in such a way that simplex σ may be reduced only by the lower node controlling it.

6.2.1 Elementary coreduction

On each node Algorithm 6 is executed to find and delete elementary coreduction pairs.

Algorithm 6 elementaryCoreduction

1. if there exists a simplex $\sigma \in$ `Simp` s.t.
   `card`$\{\tau \in \sigma.$`faces()` $|$ `not` $\tau.$`deleted`$\} = 1$
   and `canDelete`(σ), then proceed, otherwise `return false`.
2. $\sigma.$`locked` $:=$ `true`.
3. if there exists a simplex $\tau \in \sigma.$`faces()` s.t. `canDelete`(τ) then
   (a) `delete`(τ).
   (b) `delete`(σ).
   (c) `return true`.
4. otherwise
   (a) $s' :=$ the other node controlling σ.
   (b) $\tau :=$ the unique simplex in $\sigma.$`faces()` s.t. `not` $\tau.$`deleted`.
   (c) if $s'.$`canDelete`(τ) (see Lemma 2) then
       i. $s'.$`delete`(τ).
       ii. `delete`(σ).
       iii. `return true`.
   (d) else $\sigma.$`locked` $:=$ `false`.
5. `return false`.

Algorithm 7 elementaryReduction

1. if there exists a simplex $\tau \in$ `Simp` s.t.
   `card`$\{\sigma \in \tau.$`cofaces()` $|$ `not` $\sigma.$`deleted`$\} = 1$
   and `canDelete`(τ), then proceed, otherwise `return false`.
2. $\tau.$`locked` $:=$ `true`.
3. $\sigma :=$ the unique simplex in $\tau.$`cofaces()` s.t. `not` $\sigma.$`deleted`.
4. if `canDelete`(σ)
   (a) `delete`(τ).
   (b) `delete`(σ).
   (c) `return true`.
5. else if `not` $\sigma.$`locked`.
   (a) $s' :=$ the other node controlling σ.
   (b) if $s'.$`canDelete`(σ) then
       i. $s'.$`delete`(σ).
       ii. `delete`(τ).
       iii. `return true`.
   (c) else $\tau.$`locked` $:=$ `false`.
6. `return false`.

6.2.2 Elementary reduction

On each node Algorithm 7 is executed to find and delete an S-reduction pair.

6.2.3 Parallel reduction algorithm

Algorithm 8 is executed in a loop on every node. The algorithm terminates when no node in the network is able to find an S-reduction pair. The stop criterion of Algorithm 8 requires global information from the network. In consequence, the criterion cannot be implemented on the basis of purely local information as in the case of the

---

Algorithm 8 reductionAlgorithm

1. Run `createLocalFlagComplex` followed by `removeFence`.
2. Run the following code as long as there exists a node in the whole sensor network which returns `true` from `elementaryCoreduction` or `elementaryReduction` algorithm.
   (a) Run the `elementaryCoreduction` algorithm as long as there exists a node that returns `true`.
   (b) Run the `elementaryReduction` algorithm as long as there exists a node that returns `true`.

---

other algorithms presented in the paper. However, the stop criterion may be easily implemented via broadcasts.

7 Correctness

In this section we prove that the algorithm presented in the previous section correctly reduces the flag complex in the sense that the homology of the original flag complex considered as an S-complex and the homology of the reduced S-complex are the same.

Let $\mathcal{A}$ be the set of all S-reduction pairs reduced by Algorithm 8 in all nodes and let

$$\mathcal{A}^r := \{\alpha \in \mathcal{A} \mid \alpha \text{ is removed as a reduction pair}\},$$
$$\mathcal{A}^c := \{\alpha \in \mathcal{A} \mid \alpha \text{ is removed as a coreduction pair}\}.$$

For $\mathcal{A}' \subset \mathcal{A}$ set $|\mathcal{A}'| := \bigcup_{\alpha \in \mathcal{A}'} |\alpha|$, and let

$$\mathcal{K} := \mathcal{F} \setminus \mathcal{C}, \tag{6}$$
$$\mathcal{G}^r := |\mathcal{A}^r|, \tag{7}$$
$$\mathcal{G}^c := |\mathcal{A}^c|, \tag{8}$$
$$\mathcal{K}^f := \mathcal{K} \setminus \mathcal{G}^r \setminus \mathcal{G}^c. \tag{9}$$

In other words, $\mathcal{K} = \mathcal{F} \setminus \mathcal{C}$ is the S-complex resulting from removing the fence cycle from the flag complex of the sensor network, $\mathcal{G}^r$ is the subset of $\mathcal{K}$ consisting of generators removed from $\mathcal{K}$ in elementary reductions, $\mathcal{G}^c$ is the subset of $\mathcal{K}$ consisting of generators removed from $\mathcal{K}$ in elementary coreductions and $\mathcal{K}^f$ is the S-complex resulting from $\mathcal{K}$ after applying all S-reductions reduced by Algorithm 8 in all nodes.

Theorem 6 *Given a network communication graph with simple cycle $\mathcal{C}$, Algorithm 8 terminates in every node. When all copies of the algorithm complete in all nodes, the S-complex $\mathcal{K}^f$ satisfies*

$$H(\mathcal{F}, \mathcal{C}) \cong H(\mathcal{K}^f).$$

Proof Since the number of nodes is finite, also the number of all possible reductions is finite. It is clear, that in Algorithm 6 and Algorithm 7 either an S-reduction pair is reduced, or `false` is returned. Consequently either some S-reduction pair is reduced in the sensor network, or, according to point (2) of Algorithm 8, the algorithm terminates in the whole sensor network. Therefore the algorithm must terminate. Observe:

$$H(\mathcal{F}, \mathcal{C}) \cong H(\mathcal{F}/\mathcal{C}) \cong H(\mathcal{F} \setminus \mathcal{C}) = H(\mathcal{K}),$$

where the first isomorphism uses excision and the second uses Theorem 3. We need to prove that

$$H(\mathcal{K}) \cong H(\mathcal{K}^f). \tag{10}$$

For $\alpha = (\tau, \sigma) \in \mathcal{A}$ set

$$\mathcal{A}(\alpha) := \begin{cases} \{\bar{\alpha} \in \mathcal{A} \mid (\mathrm{bd}_{\mathcal{K}} \sigma \setminus \tau) \cap |\bar{\alpha}| \neq \emptyset\} & \text{if } \alpha \in \mathcal{A}^c, \\ \{\bar{\alpha} \in \mathcal{A} \mid (\mathrm{cbd}_{\mathcal{K}} \tau \setminus \sigma) \cap |\bar{\alpha}| \neq \emptyset\} & \text{if } \alpha \in \mathcal{A}^r. \end{cases}$$

In the course of running the reduction calls throughout the network a function $\lambda : \mathcal{A} \to \mathbb{N}$ is defined recursively as follows. The value $\lambda(\alpha)$ is 1 for $\alpha \in \mathcal{A}$ such that $\mathcal{A}(\alpha) = \emptyset$. Such pairs are S-reduction pairs in the original complex, form a free collection and may be reduced immediately. However, if $\mathcal{A}(\alpha) \neq \emptyset$, then the reduction of α may be performed only after all elements in $\mathcal{A}(\alpha)$ have already been reduced. This, in particular, means that the value of λ is already assigned to all elements of $\mathcal{A}(\alpha)$. Therefore, for any $\alpha \in \mathcal{A}$ we may set

$$\lambda(\alpha) := 1 + \max \{\lambda(\bar{\alpha}) \mid \bar{\alpha} \in \mathcal{A}(\alpha)\}.$$

Now, define

$$\mathcal{A}_n := \{\alpha \in \mathcal{A} \mid \lambda(\alpha) = n\}, \quad \mathcal{A}_n^r := \mathcal{A}^r \cap \mathcal{A}_n, \quad \mathcal{A}_n^c := \mathcal{A}^c \cap \mathcal{A}_n,$$
$$\mathcal{K}^0 := \mathcal{K}, \quad \mathcal{K}^n := \mathcal{K}^{n-1} \setminus |\mathcal{A}_n|.$$

We claim that for each $n \in \mathbb{N}$

(i) $\mathcal{A}_n$ is a free collection of S-reduction pairs in $\mathcal{K}^{n-1}$,
(ii) $\mathcal{K}^n$ is an S-complex.

We will prove both properties by induction on n. Since $\mathcal{A}_0 = \emptyset$ and $\mathcal{K}^0 = \mathcal{K}$, both (i) and (ii) are obvious for $n = 0$. Thus assume that the properties are satisfied for all $n < k$. To prove that property (i) holds for $n = k$ take $\alpha = (\tau, \sigma) \in \mathcal{A}_k$. Then, $\lambda(\alpha) = k$. It follows from the definition of λ that $\mathcal{A}(\alpha) \subset \mathcal{A}_1 \cup \mathcal{A}_2 \cup \cdots \cup \mathcal{A}_{k-1}$. Therefore, we have

$$\mathrm{bd}_{\mathcal{K}} \sigma \subset \{\tau\} \cup |\mathcal{A}_1 \cup \mathcal{A}_2 \cup \cdots \cup \mathcal{A}_{k-1}| \quad \text{if } (\tau, \sigma) \in \mathcal{A}_k^c, \tag{11}$$
$$\mathrm{cbd}_{\mathcal{K}} \tau \subset \{\sigma\} \cup |\mathcal{A}_1 \cup \mathcal{A}_2 \cup \cdots \cup \mathcal{A}_{k-1}| \quad \text{if } (\tau, \sigma) \in \mathcal{A}_k^r. \tag{12}$$

Algorithm 9 checkIfBoundaryless

1. for every $\sigma \in$ Simp
   (a) if σ.deleted == false then
       i. for every $\tau \in \sigma$.faces()
          if τ.deleted == false then return false;
2. return true;

Thus, from Proposition 1 and (11–12) we get

$$\mathrm{bd}_{\mathcal{K}^{k-1}}\,\sigma = \mathrm{bd}_{\mathcal{K}}\,\sigma \cap \mathcal{K}^{k-1} = \{\tau\},$$
$$\mathrm{cbd}_{\mathcal{K}^{k-1}}\,\tau = \mathrm{cbd}_{\mathcal{K}}\,\sigma \cap \mathcal{K}^{k-1} = \{\sigma\},$$

which proves that $\alpha = (\sigma, \tau)$ is and S-reduction pair in $\mathcal{K}^{k-1}$. Thus (i) holds for $n = k$ and we obtain (ii) for $n = k$ as an immediate consequence of Theorem 4(ii).

Now, let φ_n be any sequence ordering the S-reduction pairs in $\mathcal{A}_n$. Since by (i) $\mathcal{A}_n$ is free, it follows from Theorem 4(iii) that we have a well defined isomorphism $I_{\varphi_n} : H(\mathcal{K}^{n-1}) \to H(\mathcal{K}^n)$ and the composition of all the isomorphism I_{φ_n} gives the isomorphism required in (10). □

To check if the S-complex resulting from the reduction process is boundaryless, one can use Algorithm 9. One easily verifies that if the resulting S-complex is boundaryless, then all sensors return a value of true from this algorithm.

In case Conjecture 1 does not hold and the final complex is not boundaryless, we expect it to be very small so that we can easily transfer it to a single selected sensor via the communication channels available in the network. For this, we first create a spanning tree of the network using standard techniques for a distributed setting as in [2]. This may be achieved in $O(V^{1.6} + E)$, where V is a number of sensors, E is a number of edges in the flag complex. Then, we move the information along the spanning tree to its root and use the root sensor to compute the generators of the homology of the remaining complex, for instance using homology algorithm available in [31] or [32]. Finally, we use the spanning tree again to send the homology generators back to the respective simplices.

8 Verifying coverage

In this section we present an algorithm verifying the assumptions of Theorem 1. For the sake of simplicity, we assume that Conjecture 1 holds and that after applying Algorithm 8 there is left only one boundaryless 2-simplex ω. Let $\hat{\omega}$ denote the associated elementary chain which sends ω to one and everything else to zero. The nodes need to verify whether there exists a homology class $[c] \in H(\mathcal{F}, \mathcal{C})$ whose boundary is nonzero. By applying Algorithm 8 the nodes know the homology of $H(\mathcal{F}, \mathcal{C})$ but only via the isomorphic homology of $\mathcal{K}^f$, where $\mathcal{K}^f$ is given by (9). Therefore, it is necessary to find the isomorphic counterpart of $[\hat{\omega}]$ in $H(\mathcal{F}, \mathcal{C})$. To achieve this it is sufficient

to apply formula (2) repeatedly in proper order, for every S-reduction pair reduced by Algorithm 8. The respective algorithm is Algorithm 10. We make the following two assumptions concerning this algorithm:

– To store the resulting chain an extra field `coef` is added to the data structure `Simplex` and the field is initially set to `undefined`.
– The nodes remember not only simplices which were reduced but also whether they were reduced in an elementary reduction or coreduction.

In the sequel we use the natural pairing of the elements of $\mathcal{G}^r$ and $\mathcal{G}^c$ coming from S-reduction pairs in the form of a bijection

$$\mathcal{G}^r \cup \mathcal{G}^c \ni \sigma \mapsto \sigma^* \in \mathcal{G}^r \cup \mathcal{G}^c$$

which sends an element $\sigma \in \mathcal{G}^r \cup \mathcal{G}^c$ to its companion in the respective S-reduction pair.

Algorithm 10 verifyCoverage

1. for each $\sigma \in$ `Simp`[2]
   (a) if $\sigma \in \mathcal{K}^f$ set $\sigma.$`coef` := 1;
   (b) if $\sigma \in \mathcal{G}^r$ set $\sigma.$`coef` := 0;
   (c) if $\sigma \in \mathcal{G}^c$ and dim $\sigma^* \neq 1$ set $\sigma.$`coef` := 0;
2. Let $L := \{\tau \in$ `Simp`[1] $\cap \mathcal{G}^c \mid$ dim $\tau^* = 2\}$;
3. while ($L \neq \emptyset$)
   (a) for each $\tau \in L$
      i. if $\sigma.$`coef` is set for each $\sigma \in$ cbd $\tau \setminus \tau^*$
         A. $s := 0$;
         B. for each $\sigma \in$ cbd $\tau \setminus \tau^*$ do
               $s +=\sigma.$`coef` $/\kappa(\sigma, \tau)$;
         C. $\tau^*.$`coef` := s;
         D. $L := L \setminus \tau$;
4. for each $\tau \in$ `Simp`[1] such that τ is a fence edge
   (a) $s := 0$;
   (b) for each $\sigma \in$ cbd τ do $s +=\sigma.$`coef` $*\kappa(\sigma, \tau)$;
   (c) if $s \neq 0$ report "coverage verified" and exit;

The following proposition is straightforward.

Proposition 2 *If for some $\sigma \in$ Simp[2] the value of $\sigma.$`coef` is not set after completing loop (1) of Algorithm 10, then $\sigma \in \mathcal{G}^c$ and dim $\sigma^* = 1$.*

Lemma 4 *Assume that given an $n \in \mathbb{N}$ Algorithm 10 sets the value $\sigma.$`coef` for every $\sigma \in$ Simp[2] such that $(\sigma^*, \sigma) \in \mathcal{G}^c$ and $\lambda(\sigma^*, \sigma) > n$. Then $\sigma.$`coef` is set for every $\sigma \in$ Simp[2] such that $(\sigma^*, \sigma) \in \mathcal{G}^c$ and $\lambda(\sigma^*, \sigma) = n$.*

Proof Assume $\sigma_0 \in$ `Simp`[2] is such that

$$(\sigma_0^*, \sigma_0) \in \mathcal{G}^c \quad \text{and} \quad \lambda(\sigma_0^*, \sigma_0) = n.$$

Then $\tau_0 := \sigma_0^* \in L$. Let $\sigma \in \text{cbd } \tau_0 \setminus \sigma_0$. If $\sigma \in \mathcal{G}^r \cup \mathcal{K}^f$ or $\sigma \in \mathcal{G}^c$ and $\dim \sigma^* \neq 1$, then by Proposition 2 σ. coef is set already in loop (1). Consider the other case when $\sigma \in \mathcal{G}^c$ and $\dim \sigma^* = 1$. Then $\text{bd } \sigma \setminus \sigma^* \subset |\mathcal{A}(\sigma^*, \sigma)|$. Since $\tau_0 \in \text{bd } \sigma$ and $\tau_0 \neq \sigma^*$, we get $(\sigma_0^*, \sigma_0) \in \mathcal{A}(\sigma^*, \sigma)$. Therefore $\lambda(\sigma^*, \sigma) > \lambda(\sigma_0^*, \sigma_0) = n$ and by the assumption σ_0. coef is set also in this case. It follows that when τ_0 is considered in line $(3.a)$, the condition in line $(3.a.i)$ is satisfied and consequently σ. coef $= \tau_0^*$. coef is set in line $(3.a.i.C)$. $\square$

Theorem 7 *Algorithm 10 terminates in all nodes. Moreover, the assumptions of Theorem 1 are satisfied if and only if at least one node exits the algorithm with the report "coverage verified".*

Proof First we will show that the algorithm sets σ. coef for each $\sigma \in \text{Simp}[2]$. By Proposition 2 we need to consider only the case when $\sigma \in \mathcal{G}^c$ and $\dim \sigma^* = 1$. However, this case follows immediately from Lemma 4 by induction with respect to $n := N - \lambda(\sigma^*, \sigma)$, where

$$N := \max \{ \lambda(\alpha) \mid \alpha \in \mathcal{A} \}.$$

Observe that whenever σ. coef is set in line $(3.a.i.C)$, then σ^* is removed from L in line $(3.a.i.D)$. Therefore, the loop in line (3) completes and consequently also Algorithm 10 completes in every node.

It is now straightforward to verify that after completing loop (3), the fields σ. coef for $\sigma \in \text{Simp}[2]$ store a chain $c_\omega \in \mathcal{K}$ whose homology is the image of $[\hat{\omega}]$ under the isomorphism established in Theorem 6.

Now, the variable s evaluated in loop (4) stores $\langle \partial c_\omega, \tau \rangle$ for some $1-$simplex τ of the fence. Therefore, if at least one node reports "coverage verified", then the assumptions of Theorem 1 are satisfied and the coverage is indeed archived. If the assumptions are not satisfied, then obviously no node can report "coverage verified." $\square$

In case, when there is more then one boundaryless 2-simplex left at the end of the reduction process, the described algorithm may be called for each of the remaining 2-simplices and if at least one of the nodes reports "coverage verified" for at least one of the remaining 2-simplices, the assumptions of Theorem 1 are satisfied. The adaptation of Algorithm 10 to the case when Conjecture 1 does not hold is only slightly more complicated. It requires finding the homology generators of $H(\mathcal{K}^f)$ by some algebraic means, for instance by applying the algorithm described in [1] and then starting Algorithm 10 with the chains representing the homology generators of $H(\mathcal{K}^f)$ instead of the elementary chain $\hat{\omega}$. Therefore, the extra assumptions from the beginning of the section are not restrictive.

9 On complexity

The full complexity analysis of the presented algorithms is not possible; in this paper we consider neither a concrete communication model nor the synchronization methods needed to collate algorithm phases. However, even at the level of generality with

which our analysis is performed, some basic remarks toward a more complete analysis are possible. The details of a full complexity analysis, given length and technicality, are to be published elsewhere.

We make the following simplifying assumptions:

1. Every sensor can send an integer to its neighbors without error and in one time unit.
2. A global synchronization method which may be used to synchronize in constant time the consecutive steps of the algorithms in the sensors is available.
3. The implementation of the algorithm is based on best performing data structures for finding simplices, queues for reduction candidates etc.

One can show that under these assumptions the algorithm constructing the flag complex (Algorithm 2) and the algorithm reducing it (Algorithm 8) may be implemented in such a way that the complexity for one sensor (compare network time defined in Sect. 10) is $O(n^2 + K(n \log K + n^2))$, where:

1. n is the upper bound for the number of sensors in the neighborhood of a given sensor.
2. K is the upper bound for the number of simplices which have the given sensor as a vertex.

In other words, the complexity depends only on the local size of the network, and not on the size of the whole network as in the case of centralized computations. This justifies the utility of the presented algorithms in the context of large-scale sensor networks of bounded density. The analysis depends on the Conjecture 1, because the estimate does not include the potential postprocessing of the final S-complex which is not boundaryless, although we expect this cost to be negligible even if nonzero.

10 Simulations

The parallel reduction algorithms presented in the paper have been implemented in Java programming language [33]. The code executes the algorithms on each node in a separate thread, so that simulations may be performed with only a few or even one processor unit. The real time of the network is also simulated.

In this section we present the experimental results showing the advantages of the distributed homology computation. In the presented series of experiments, the following configuration, referred to as the *base test* will be used. The base test is a small network of 23 nodes randomly placed on a 4×4 units rectangle with the communication radius fixed at 2 units. This base test will be copied and shifted and new fence cycles created. This process is performed to obtain a reasonably uniform distribution of the nodes on larger networks and to avoid excessive local clustering which results in high dimensional flag complexes.

We first define a family L^n of networks on a sequence of 'linear' domains, $L^n = \{L_i\}_{i=0}^n$, where L_i is defined recursively as follows:

1. L_0 is the base test case with a rectangular fence cycle $\mathcal{C}$ around the network; and
2. L_i is created from L_{i-1} by placing a new copy of the base test on the right side of L_{i-1} with a new rectangular fence cycle $\mathcal{C}$ around the network.

Fig. 5 Network Time and Running Time for the family L^{53} of networks

Another family $S^n = \{ S_i \}_{i=0}^n$ of 'square' networks is defined recursively as follows:

1. S_0 is the base test case with a rectangular fence cycle $\mathcal{C}$ around the network; and
2. S_i is created from S_{i-1} by placing new copies of the base test: i-copies on the right, i-copies on the bottom side, and one on the bottom-right. A new rectangular fence $\mathcal{C}$ around the network is created.

Let C_s denote the CPU time used by node s during the simulation and let R_s be the real time used by node s, i.e., the world time passing from the moment the node starts the computations until the moment it completes. We define the *network time* as $\max_{s \in S} R_s$ and the *running time* as $\sum_{s \in S} C_s$. In other words, the network time is the time required by the real network to do the computation. In the case of simulations performed on a single machine with a single CPU, the running time is the CPU usage for the whole simulation. We define the network efficiency as the ratio

$$\texttt{Network Efficiency} := \frac{\texttt{Network Time}}{\texttt{Running Time}} \cdot 100\,\%$$

For the families L^{53} and S^{10} we ran simulations and measured network time, running time and network efficiency. Figure 5 presents the network and running time as the functions of the size of the network for the family L^{53}. The network efficiency for the same data is presented in Fig. 6. Analogous results for the family S^{10} are in Fig. 7 and Fig. 8. The outcome of the experiments clearly indicates the advantage gained through distributed computation.
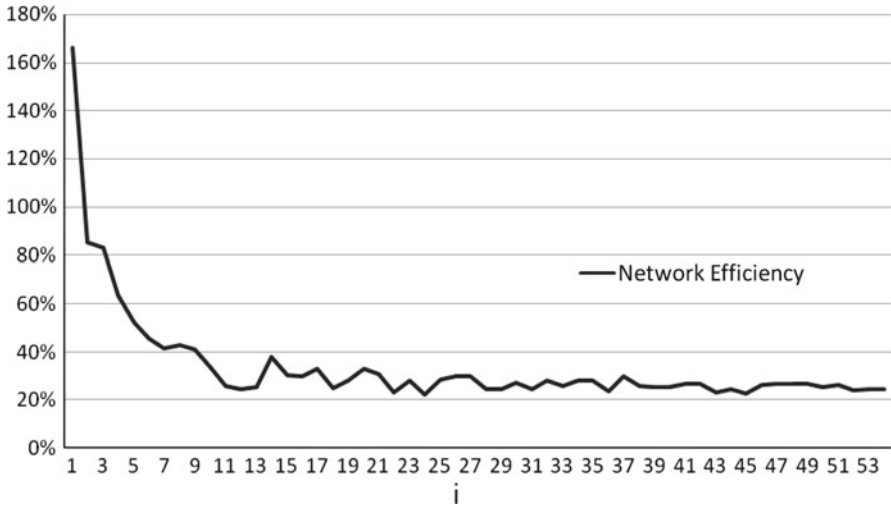
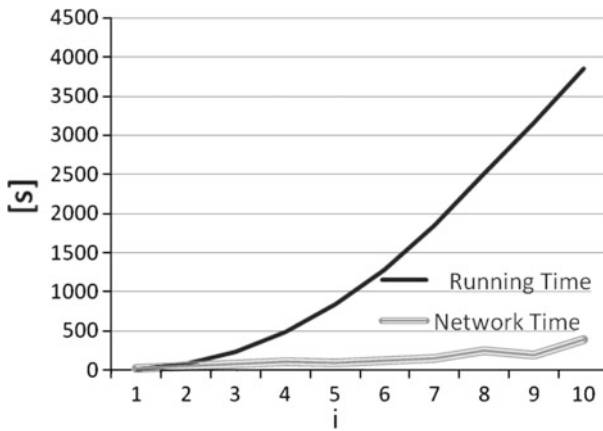**Fig. 6** Efficiency of the network for the family L^{53} of networks

**Fig. 7** Network Time and Running Time for the family S^{10} of networks

Fig. 8 Efficiency of the network for the family S^{10} of networks

Appendix: Primer on homology theory

This abbreviated introduction to homological tools assumes a knowledge of basics from algebra and topology. From algebra, the notions of rings (algebraic generalizations of $\mathbb{Z}$) and modules (algebraic generalizations of vector spaces) are assumed, as are homomorphisms (algebraic generalizations of linear transformations), kernels, and the like. From topology, only basic notions of simplicial complexes are needed.

Chain complexes

Homology counts objects with cancellation to provide a topological invariant in algebraic terms. The simplest version of homology is *simplicial homology* of a simplicial complex. Fix a finite simplicial complex X. The building blocks of a rudimentary homology for X are as follows.

1. Fix a coefficient ring R.
2. Grade the simplices of X by dimension.
3. Define *q-chains C_q* as the R-module with basis the (oriented) q-simplices of X.
4. Consider the *boundary maps*—the linear transformations $\partial : C_q \to C_{q-1}$ which send a basis q-simplex to its boundary faces (as an abstract sum of basis $(q-1)$-simplices, each with coefficient ± 1 depending on orientation).

The sequence of chains and boundary maps are assembled into a *chain complex*—a sequence $\mathcal{C}$ of R-modules C_q and homomorphism $\partial_q : C_q \to C_{q-1}$ with $\partial_q \circ \partial_{q+1} = 0$ for all q. A chain complex may be written out as a diagram,

$$\cdots \longrightarrow C_q \xrightarrow{\partial_q} C_{q-1} \xrightarrow{\partial_{q-1}} \cdots \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0 . \qquad (13)$$

or as a single object $\mathcal{C} = (C_*, \partial_*)$ and to write ∂ for the boundary operator acting on any chain of unspecified grading. Chain complexes need not be generated by simplices of a simplicial complex: they are decidedly algebraic devices. The key requirement for a chain complex is that the *boundary of a boundary is null*: $\partial_q \circ \partial_{q+1} = 0$ for all q.

Homology

Homology counts equivalence classes of certain chains with regards to the boundary maps. A *cycle* of $\mathcal{C}$ is a chain with empty boundary, i.e., an element of ker ∂. Homology is an equivalence relation on cycles of $\mathcal{C}$. Two cycles in $Z_q := \ker \partial_q$ are *homologous* if they differ by an element of $B_q := \operatorname{im} \partial_{q+1}$. The *homology* of $\mathcal{C}$ is the sequence of

quotient modules $H_q(\mathcal{C})$, for $q \in \mathbb{N}$, given by:

$$H_q(\mathcal{C}) := Z_q/B_q = \ker \partial_q \big/ \operatorname{im} \partial_{q+1}. \tag{14}$$

Elements of $H_q(C)$ are *homology classes*. We write $H(\mathcal{C})$ to denote the full module of graded homologies $H_q(\mathcal{C})$. When $\mathcal{C}$ is the chain complex generated by a simplicial complex X, we write $H(X)$ for its homology: it is an invariant of the space X up to homotopy.

Relative homology

Homology is defined for any chain complex $\mathcal{C} = (C_q, \partial_q)$. Although taking C_q to be a module generated by q-dimensional simplices is common, it is by no means exclusive. Many of the constructs of this paper rely on a modified chain complex giving rise to *reduced homology* relative to a subcomplex. Let $A \subset X$ be a (necessarily closed) subcomplex of a simplicial complex X. Then both X and A define simplicial chain complexes, with the result that $C_q(A) \subset C_q(X)$ is a submodule for all q. Thus, one can take the quotient module $C_q(X, A) := C_q(X)/C_q(A)$, so that a relative chain is an equivalence class of chains relative to simplices in A. The boundary map ∂_q extends naturally to $\partial_q : C_q(X, A) \to C_{q-1}(X, A)$ in a manner that preserves the $\partial^2 = 0$ condition. Therefore, the *relative homology* $H_q(X, A)$ is well-defined and measures *relative cycles* (chains in X whose boundaries lie in A) modulo relative boundaries.

Functoriality

A *chain map* is a map $\varphi : \mathcal{C} \to \mathcal{C}'$ between chain complexes that is a homomorphism on chains respecting the grading and commuting with the boundary maps. This is best expressed in the form of a *commutative diagram*:

$$
\begin{array}{ccccccccc}
\cdots & \longrightarrow & C_{q+1} & \xrightarrow{\ \partial\ } & C_q & \xrightarrow{\ \partial\ } & C_{q-1} & \longrightarrow & \cdots \\
& & \downarrow{\varphi} & & \downarrow{\varphi} & & \downarrow{\varphi} & & \\
\cdots & \longrightarrow & C'_{q+1} & \xrightarrow{\ \partial'\ } & C'_q & \xrightarrow{\ \partial'\ } & C'_{q-1} & \longrightarrow & \cdots
\end{array}
\tag{15}
$$

Commutativity means that homomorphisms are path-independent in the diagram; e.g., $\varphi \circ \partial = \partial' \circ \varphi$. Chain maps are the analogues of continuous maps, given their respect for the boundary operators: neighbors are sent to neighbors.

A chain map φ induces homomorphisms on homology groups, written $H(\varphi) : H(\mathcal{C}) \to H(\mathcal{C}')$, sending $[\zeta] \in H_q(\mathcal{C})$ to $[\varphi(\zeta)] \in H_q(\mathcal{C}')$. The reader may check that this is a well-defined homomorphism. We denote by $H(\varphi)$ the full sequence of *induced homomorphisms* on homology. Homology is *functorial*, meaning that induced homomorphisms respect composition of chain maps. Specifically,

1. The identity chain map induces the identity isomorphism on homology, id: $H(\mathcal{C}) \to H(\mathcal{C})$.
2. Composable chain maps φ and ψ satisfy $H(\psi \circ \varphi) = H(\psi) \circ H(\varphi)$.

Exact sequences

Homology computations are greatly aided by a theoretical tool called an *exact sequence*. Any chain complex $\mathcal{C} = (C_q, \phi_q)$ of R-modules and homomorphism is *exact* when its homology vanishes: $\ker \phi_q = \operatorname{im} \phi_{q+1}$ for all q. An exact chain complex is the chain analogue of a nullhomologous space. Exact sequences are most often used to prove isomorphisms between various homologies by means of zeroing out terms in an exact sequence. For example, if some subsequence of an exact sequence reads as:

$$\cdots \longrightarrow 0 \longrightarrow G \xrightarrow{\phi} H \longrightarrow 0 \longrightarrow \cdots$$

then it follows that $\phi : G \to H$ is an isomorphism. More generally, the kernel and cokernel of a homomorphism $\phi : G \to H$ fit into an *short exact sequence*:

$$0 \longrightarrow \ker \phi \longrightarrow G \xrightarrow{\phi} H \longrightarrow \operatorname{coker} \phi \longrightarrow 0$$

The most important examples of exact sequences are those relating homologies of various spaces and subspaces. These almost always derive from the following result in homological algebra:

Theorem 8 (Snake Lemma) *Any short exact sequence of chain complexes*

$$0 \longrightarrow \mathcal{A} \xrightarrow{i} \mathcal{B} \xrightarrow{j} \mathcal{C} \longrightarrow 0$$

induces the long exact sequence:

$$\longrightarrow H_q(\mathcal{A}) \xrightarrow{H(i)} H_q(\mathcal{B}) \xrightarrow{H(j)} H_q(\mathcal{C}) \xrightarrow{\delta} H_{q-1}(\mathcal{A}) \xrightarrow{H(i)} . \quad (16)$$

Moreover, the long exact sequence is functorial: a commutative diagram of short exact sequences and chain maps

$$
\begin{array}{ccccccccc}
0 & \longrightarrow & \mathcal{A} & \longrightarrow & \mathcal{B} & \longrightarrow & \mathcal{C} & \longrightarrow & 0 \\
& & \downarrow{f} & & \downarrow{g} & & \downarrow{h} & & \\
0 & \longrightarrow & \tilde{\mathcal{A}} & \longrightarrow & \tilde{\mathcal{B}} & \longrightarrow & \tilde{\mathcal{C}} & \longrightarrow & 0
\end{array}
$$

induces a commutative diagram of long exact sequences

$$
\begin{array}{ccccccccc}
\longrightarrow & H_q(\mathcal{A}) & \longrightarrow & H_q(\mathcal{B}) & \longrightarrow & H_q(\mathcal{C}) & \xrightarrow{\delta} & H_{q-1}(\mathcal{A}) & \longrightarrow & . \quad (17) \\
& \downarrow{H(f)} & & \downarrow{H(g)} & & \downarrow{H(h)} & & \downarrow{H(f)} & & \\
\longrightarrow & H_q(\tilde{\mathcal{A}}) & \longrightarrow & H_q(\tilde{\mathcal{B}}) & \longrightarrow & H_q(\tilde{\mathcal{C}}) & \xrightarrow{\delta} & H_{q-1}(\tilde{\mathcal{A}}) & \longrightarrow &
\end{array}
$$

An *exact sequence of chain complexes* means that there is a short exact sequence in each grading, and these short exact sequences fit into a commutative diagram with respect to the boundary operators. The induced *connecting homomorphism* $\delta : H_q(\mathcal{C}) \to H_{q-1}(\mathcal{A})$ comes from the boundary map in $\mathcal{C}$.

Given $A \subset X$ a subcomplex, the following short sequence is exact:

$$0 \longrightarrow \mathcal{C}(A) \xrightarrow{i} \mathcal{C}(X) \xrightarrow{j} \mathcal{C}(X, A) \longrightarrow 0 \, ,$$

where $i : A \hookrightarrow X$ is inclusion and $j : (X, \varnothing) \hookrightarrow (X, A)$ is an inclusion of pairs. This yields the *long exact sequence of the pair* (X, A):

$$\longrightarrow H_q(A) \xrightarrow{H(i)} H_q(X) \xrightarrow{H(j)} H_q(X, A) \xrightarrow{\delta} H_{q-1}(A) \longrightarrow \, . \quad (18)$$

The connecting homomorphism δ takes a relative homology class $[\alpha] \in H_q(X, A)$ to the homology class $[\partial \alpha] \in H_{q-1}(A)$.

References

1. Arai, Z., Hayashi, K., Hiraoka, Y.: Mayer-Vietoris sequences and coverage problems in sensor networks. Jpn. J. Ind. Appl. Math. **28**, 237–250 (2011)
2. Awerbuch, B., Gallager, R.: A new distributed algorithm to find breadth first search trees. IEEE Trans. Inf. Theory. **33**(3), 315–322 (1987)
3. Barrière, L., Fraigniaud, P., Narayanan, L.: Robust position-based routing in wireless ad hoc networks with unstable transmission ranges. In: Proc. Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications. (2001)
4. Carlsson, G., de Silva, V.: Zigzag persistence. In: Proc. Found. of Computational Mathematics. (2009)
5. Carlsson, G., de Silva, V., Morozov, D.: Zigzag persistent homology and real-valued functions. In: Proc. Symp. on Comput. Geometry. (2009)
6. Chambers, E., de Silva, V., Erickson, J., Ghrist, R.: Rips complexes of planar point sets. Discret. Comput. Geom. **44**(1), 75–90 (2010)
7. Cortes, J., Martinez, S., Karatas, T., Bullo, F.: Coverage control for mobile sensing networks. In: Proc. IEEE Int. Conf. Robot. Autom., vol. 2, pp. 1327–1332. Washington, DC (2002)
8. Damian, M., Pandit, S., Pemmaraju, S.: Local approximation schemes for topology control. In: Proc. ACM Symp. on Prin. of Dist. Comput. (PODC), pp. 208–217 (2006)
9. de Silva, V., Ghrist, R.: Coordinate-free coverage in sensor networks with controlled boundaries via homology. Int. J. Robot. Res. **25**, 1205–1222 (2006)
10. de Silva, V., Ghrist, R.: Homological sensor networks. Notices Am. Math. Soc. **54**(1), 10–17 (2007)
11. Eckmann, B.: Harmonische funktionen und randwertaufgaben einem komplex. Comment. Math. Helvetici. **17**, 240–245 (1945)
12. Estrin, D., Culler, D., Pister, K., Sukhatme, G.: Connecting the physical world with pervasive networks. IEEE Pervasive Comput. **1**(1), 59–69 (2002)
13. Fekete, S., Kröller, A., Pfisterer, D., Fischer, S.: Deterministic boundary recongnition and topology extraction for large sensor networks. In: Algorithmic Aspects of Large and Complex Networks. (2006)
14. Gelfand, S., Manin, Y.: Methods of Homological Algebra, 2nd edn. Springer, Berlin (2003)
15. Ghrist, R., Hiraoka, Y.: Applications of sheaf cohomology and exact sequences to network coding. Proc. NOLTA: Nonlinear Theory Appl. 266–269 (2011)
16. Kempe, D., Dobra, A., Gehrke, J.: Computing aggregate information using gossip. In: Proc. Foundations of Computer Science, Cambridge, MA (2003)
17. Koskinen, H.: On the coverage of a random sensor network in a bounded domain. In: Proceedings of 16th ITC Specialist Seminar, pp. 11–18. (2004)

18. Kuhn, F., Wattenhofer, R., Zollinger, A.: Ad-hoc networks beyond unit disk graphs. Wirel. Netw. **14**(5), 715–729 (2008)
19. Li, X.-Y., Wan, P.-J., Frieder, O.: Coverage in wireless ad-hoc sensor networks. IEEE Trans. Comput. **52**(6), 753–763 (2003)
20. Meguerdichian, S., Koushanfar, F., Potkonjak, M., Srivastava, M.: Coverage problems in wireless ad-hoc sensor networks. In: IEEE INFOCOM, pp. 1380–1387 (2001)
21. Mrozek, M., Batko, B.: Coreduction homology algorithm. Discret. Comput. Geom. **41**, 96–118 (2009)
22. Mrozek, M., Wanner, Th.: Coreduction homology algorithm for inclusions and persistent homology. Comput. Math. Appl. **60**(10), 2812–2833 (2010). doi:10.1016/j.camwa.2010.09.036
23. Muhammad, A., Egerstedt, M.: Control using higher order Laplacians in network topologies, In *Proc. of the 17th International Symposium on Mathematical Theory of Networks and Systems*, pp. 1024–1038. (2006)
24. Muhammad, A., Jadbabaie, A.: Decentralized computation of homology groups in networks by gossip. In: Proc. of American Control Conference, pp. 3438–3443 (2007)
25. Robinson, M.: Inverse problems in geometric graphs using internal measurements. arXiv:1008.2933v1 (2010)
26. Robinson, M.: Asynchronous logic circuits and sheaf obstructions. arXiv:1008.2729v1 (2010)
27. Tahbaz Salehi, A., Jadbabaie, A.: Distributed coverage verification in sensor networks without location information. IEEE Trans. Autom. Control **55**(8), 1837–1849 (2010)
28. Tahbaz Salehi, A., Jadbabaie, A.: Distributed coverage verification in sensor networks without location information. In: IEEE Conference on Decision and Control (2008)
29. Xue, F., Kumar, P.R.: The number of neighbors needed for connectivity of wireless networks. Wirel. Netw. **10**(2), 169–181 (2004)
30. Zhang, H., Hou, J.: Maintaining coverage and Connectivity in large sensor networks. In: International Workshop on Theoretical and Algorithmic Aspects of Sensor, Ad hoc Wireless and Peer-to-Peer Networks, Florida (2004)
31. The RedHom homology algorithms library: http://redhom.ii.uj.edu.pl
32. Computational Homology Project: http://chomp.rutgers.edu
33. Sensor Network Simulator: http://redhom.ii.uj.edu.pl/sensors/