Homological Sensor Networks

Vin de Silva and Robert Ghrist

Sensors and Sense-ability

A sensor is a device that measures some feature of a domain or environment and returns a signal from which information may be extracted. Sensors vary in scope, resolution, and ability. The information they return can be as simple as a binary flag, as with a metal detector that beeps to indicate a detection threshold being crossed. A more complex sensor, such as a video camera, can return a signal requiring sophisticated analysis to extract relevant data.

An increasingly common application for sensors is to scan a region for a particular object or substance. For example, one might wish to determine the existence and location of an outbreak of fire in a national forest. Questions of more interest to national security involve detection of radiological or biological hazards, hidden mines and munitions, or specific individuals in a crowd. All of these scenarios pose difficult and challenging data management problems.

Numerous strategies exist, aided by the fact that sensor technology provides an expansive array of available hardware. A fundamental dichotomy exists in the approach to sensing an environment based on the number and complexity of sensors. For a fixed cost (monetary, or perhaps "total complexity"), one can deploy a small number of sophisticated "global" sensors with high signal complexity and precise readings. In contrast, one can deploy a large number of small, coarse, "local" devices that may have large uncertainties in their readings. Different strategies are appropriate for different tasks. The human body contains examples of sensor systems with a small multiplicity of highly complex devices (for sight) as well as vast networks of local sensors (for touch).

Technology promises to push the envelope on both sides of this spectrum, yielding new types of powerful, global sensors, as well as local sensors of surprisingly small size. The relevant question for the mathematician is which types of mathematics will be useful in analyzing sophisticated sensor networks.

It may be that the most exciting possibilities lie in the domain of the small. Swarms of local sensors at micro- or nanoscale have the potential to revolutionize the way that we think about security and surveillance problems [4]. However, this brings with it the difficulty of integration. How does one collect local information and collate it into global environmental data?

From Local to Global

Fortunately, mathematicians have spent centuries carefully contemplating local-to-global transitions. The very term we use to indicate the collection and collation of local data—*integration*—harks back to the well established means of relating local information about a function (pointwise derivatives) with a global quantity (the integral).

A more relevant example for our purposes is to be found in simple ideas about the topology of surfaces. What are the global features of a surface given "local data" in the form of a triangulation? The Classification Theorem for Surfaces implies that the Euler characteristic $\chi(\Sigma)$ suffices to determine the homeomorphism type of a closed orientable surface Σ. The computation is as simple as one could hope for:

$$\chi(\Sigma) = \#V - \#E + \#F,$$

Vin de Silva is assistant professor of mathematics at Pomona College. His email address is `vin.desilva@pomona.edu`.

Robert Ghrist is associate professor of mathematics at the University of Illinois, Urbana-Champaign. His email address is `ghrist@math.uiuc.edu`.

The authors acknowledge the support of DARPA (Defense Advanced Research Projects Agency) and the National Science Foundation.

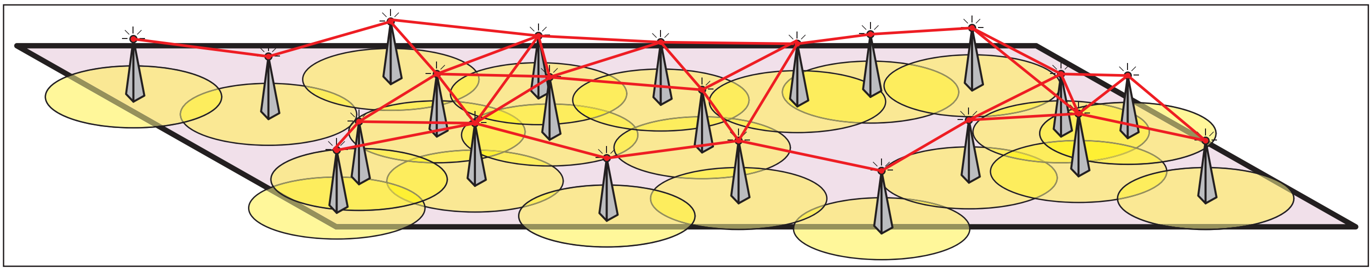

Figure 1. A network of small, local sensors samples an environment at a set of nodes. How can one answer global questions from this network of local data?

where the triangulated surface Σ has $\#V$ vertices, $\#E$ edges, and $\#F$ faces.

The information contained in $\chi(\Sigma)$ is not restricted to topological classification. The Euler characteristic $\chi(\Sigma)$ can be used to infer geometric properties of Σ (specifically, the Gauss curvature, via the Gauss-Bonnet Theorem) and dynamical properties of Σ (specifically, the number and types of fixed points of a vector field, via the Hopf Index Theorem).

The efficacy of the Euler characteristic in this example is a consequence of the restricted nature of surfaces. For a more arbitrary space, the challenge of characterizing global features of the space becomes a more fundamental problem in topology. With surfaces, simple arithmetic suffices to determine global properties. For arbitrary complexes, more sophisticated algebraic topology is required. Roughly speaking, algebraic topology provides two ways in which to associate to a given space X a collection of algebraic objects that gauge the global features of X.

The first such set of invariants are the *homotopy groups*, $\pi_k(X)$, for $k = 0, 1, \ldots$, the *fundamental group* $\pi_1(X)$ being very well known. These groups measure in how many and which ways one can map a k-dimensional sphere S^k into X, two spheres in X being deemed equivalent if they are homotopic relative to some fixed basepoint. Homotopy groups comprise very powerful data; however, they are in practice quite difficult to compute. The general computation of homotopy groups of spheres is unknown and indeed is the premier unsolved problem in algebraic topology at this time.

The second set of invariants provide a weaker but more computable option. These are the *homology groups*, $H_k(X)$, for $k = 0, 1, \ldots$. (Properly speaking, homology defers to its algebraic dual—the cohomology groups $H^k(X)$—as a finer invariant.) Instead of measuring k-spheres in a space up to homotopy, homology measures certain types of *chains*, or objects built from simple oriented pieces: *simplices*. These simplices are defined differently depending on the type of homology used. The simplest instantiation is that of a simplicial complex X, where the combinatorial simplices from which X is built form a basis for simplicial chains. The elements of $H_k(X)$ are *cycles*, or chains with vanishing boundary, and two k-cycles are deemed *homologous* if there is an oriented $(k+1)$-chain that has as its boundary the pair of cycles (with opposite orientation).

Like homotopy groups, the homology groups are an invariant of the homotopy type of the underlying space. This explains why the Euler characteristic χ of a surface is independent of both the triangulation and the homeomorphism type of the surface: χ is the alternating sum of the dimensions of the homology groups.

Unlike homotopy groups, homology groups can be computed via linear algebra. Recent advances in algorithms for the rapid computation of homology (see [7] and references therein) make this a feasible tool for realistic problems in science and engineering.

Blanket Coverage

Motivated by the potential of pervasive computing in sensor-rich environments [4], we consider a class of simple sensors that can solve global problems based on local communication.

For concreteness, we consider the case where nodes lie in a planar Euclidean domain with polygonal boundary. Each node can perform some sensing task within a certain radially symmetric neighborhood. Within this *coverage disk*, the sensor performs its unspecified task, whether it involves video surveillance, detection of radiological or biohazard material, motion detection, etc. We do not model this sensing task at all: it is completely implicit except for the assumption that it is radially symmetric. For such a network, we consider the problem of *blanket coverage*.

> Does the union of the coverage discs about the nodes cover the domain $\mathcal{D}$?

We wish to solve this problem using small-scale (and therefore cheap) devices without GPS or other sophisticated positioning systems. The intended lesson is that topological methods permit sensors that are remarkably minimal, having no means of measuring distance, orientation, or location in their environment.

The coverage problem is of clear significance to security and surveillance. A similar coverage problem vexes anyone with a cell phone in an area of

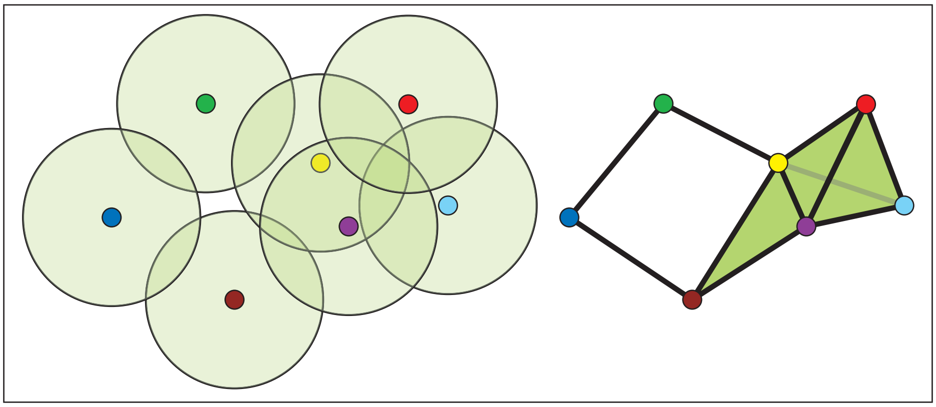

Figure 2. The Čech complex of a cover by convex sets captures the homotopy type of the cover.

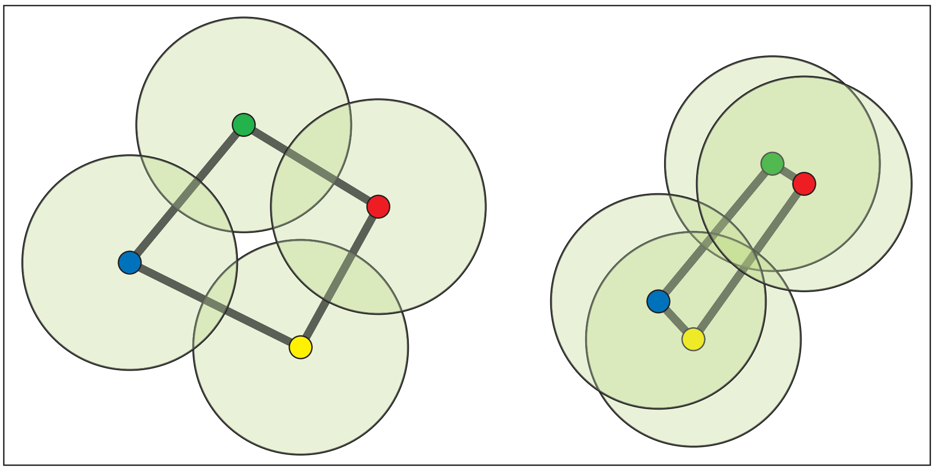

Figure 3. Changing the positions of nodes can change the topology of the radius r_c cover without changing the topology of the radius r_b network graph.

low cell phone tower density. This latter coverage problem is simpler because the network of cell phone towers is fixed and intentional. The company that built the towers knows exactly where they were built and is certain that the towers have not moved. One can thus compute the union of the coverage discs "by hand" with ease (assuming no hardware failure). Standard algorithms from computational geometry can check for holes quickly, even in cases with many nodes, so long as the node positions are known.

The scenario that we envision differs in that there is no means of determining relative position. This is not an insurmountable difficulty. Indeed, there is an extensive literature on probabilistic methods for coverage problems in networks of randomly distributed points. See, e.g., [8]. Unfortunately, these methods have very strong assumptions on the uniformity or density of the random distribution of points. We would like to solve coverage problems in more realistic settings where one "dumps a bucketful" of sensors in a field, forest, or ocean and then queries the network, perhaps after environmental influences have moved the sensors to unknown positions.

A Simple Local Network

What minimal capabilities must the sensor nodes possess, for there to be a solution (or reasonable partial solution) to the blanket coverage problem? We focus on node-to-node communication. Assume that each node broadcasts its unique ID and listens to determine its neighbors. These unique IDs may take the form of RFID tags.

The one strong assumption we make concerns the boundary of the domain $\mathcal{D}$ in which the nodes lie. We suppose (for now) that the vertices of the polygonal boundary $\partial\mathcal{D}$ are defined by special *fence nodes* in a known cyclic configuration (although their coordinates remain unknown). Our precise assumptions are as follows:

A1: Nodes $\mathcal{X}$ broadcast their unique ID numbers. Each node can detect the identity of any node within *broadcast* radius r_b.

A2: Nodes have radially symmetric covering domains of *cover* radius $r_c \geq r_b/\sqrt{3}$.

A3: Nodes $\mathcal{X}$ lie in a compact connected domain $\mathcal{D} \subset \mathbb{R}^2$ whose boundary $\partial\mathcal{D}$ is connected and piecewise-linear with vertices marked *fence nodes* $\mathcal{X}_f$.

A4: Each fence node $v \in \mathcal{X}_f$ knows the identities of its neighbors on $\partial\mathcal{D}$ and these neighbors both lie within distance r_b of v.

To summarize, each node is aware of the identities of those nodes that are within broadcast range r_b. The orientations and distances of these neighboring nodes are unknown. The fence nodes have two additional pieces of data: (1) they know that they are on the boundary of the domain; and (2) each knows the identities of the two neighboring fence nodes.

Apart from the fence nodes (which are used to simplify the statements of theorems), the type of information that this network encodes is very similar to that encoded by a simplicial complex. Local combinatorial data about how elementary pieces are assembled give rise to a global object whose large-scale topological features are revealing.

Simplices for Sensors

The obvious way to begin is to build the *network graph* of the system. This is a combinatorial graph, Γ, in which vertices correspond to the labeled nodes and (undirected) edges correspond to pairs of nodes that are in mutual broadcast range (within distance r_b). In this graph, the boundary $\partial\mathcal{D}$ is naturally identified with a particular cycle $\mathcal{F} \subset \Gamma$ traversing the fence nodes, thanks to A4. The problem at hand is to determine whether the set $\mathcal{U}$ given by the union of radius r_c balls at $\mathcal{X}$ contains the domain $\mathcal{D}$. The input for this problem is the pair of graphs $(\Gamma, \mathcal{F})$.

Determining the topology of a union of balls is a classical problem and is easily solved using the notion of a Čech complex (also known as the *nerve*). Given a collection of sets $\mathcal{U} = \{U_\alpha\}$, the *Čech complex* of $\mathcal{U}$, $C(\mathcal{U})$, is the abstract simplicial complex whose k-simplices correspond to nonempty intersections of $k+1$ distinct elements of $\mathcal{U}$. Thus, the vertices are in bijective correspondence with the cover sets U_α, and edges of $C(\mathcal{U})$ are in bijective correspondence with nonempty intersections between two cover sets. Higher order intersections generate higher dimensional simplices: see Figure 2.

Theorem 1. *[The Čech Theorem] If the sets $\{U_\alpha\}$ and all nonempty finite intersections are contractible, then the union $\bigcup_\alpha U_\alpha$ has the homotopy type of the Čech complex C.*

The equivalence in the Čech theorem is functorial, and in particular there is a relative version that gives us the following result.

Corollary 2. *Under assumptions A1-A4 above, the coverage area $\bigcup_\alpha U_\alpha$ contains the domain $\mathcal{D}$ if and only if the fence 1-cycle $\mathcal{F}$ is null-homologous in $C(\mathcal{U})$.*

This would appear to be exactly what one wants for sensor networks. Unfortunately, it is not possible to compute the Čech complex from the network graph Γ alone. Precise distances between nodes are needed to determine the higher-dimensional simplices of $C(\mathcal{U})$. All we have are two radii: the broadcast radius r_b and the coverage radius r_c. For no (physically realistic) choice of these radii can the radius r_c Čech complex be derived from the radius r_b network graph. It is not even possible to recover the homotopy type of $C(\mathcal{U})$. See Figure 3 for one example of the difficulty.

On the other hand, with the bound on coverage and broadcast radii in A2, it follows that for any triple of nodes that are in pairwise communication distance, the convex hull of these nodes in $\mathbb{R}^2$ is contained in the cover $\mathcal{U}$. The extremal case, in which all three nodes are at pairwise distance r_b, yields an equilateral triangle in $\mathbb{R}^2$ that is covered by balls at the nodes of radius r_c only if $r_c \geq r_b/\sqrt{3}$.

This motivates the following construction. We consider the network graph as the 1-dimensional skeleton of a larger simplicial complex. Denote by $\mathcal{R}$ the largest simplicial complex whose 1-skeleton is the network graph. That is, for every collection of k nodes that are pairwise within distance r_b, we assign an abstract $k-1$ simplex. This is also known as the *flag complex* associated to the network graph.

A nearly identical construction was used by Vietoris in the 1930s in the beginnings of homology theory [9]. It was largely forgotten and later reformulated by Rips in his work on geometric groups. Given a set of points $\mathcal{X} = \{x_\alpha\} \subset \mathbb{R}^n$ in Euclidean n-space and a fixed radius ϵ, the **Vietoris–Rips complex** of $\mathcal{X}$ is the abstract simplicial complex whose k-simplices correspond to unordered $(k+1)$-tuples of points in $\mathcal{X}$ that are pairwise within Euclidean distance ϵ of each other.

For brevity, we refer to the complex $\mathcal{R}$ constructed above as the **Rips complex** of the network, with the radius r_b understood implicitly. Unfortunately, the Rips complex does not necessarily capture the topology of the union of cover discs: we have traded accuracy for computability. In the remainder of this article, we will outline two methods for extracting coverage information from a Rips complex, the latter of which infers Čech data.

The Homological Criterion

The Rips complex does contain enough topological information about the cover to certify coverage, if the cover is sufficiently robust. One might guess that the right criterion measures $H_1(\mathcal{R})$, since $H_1(\mathcal{U})$ collates holes in the cover. For reasons to be seen, it is more natural to consider the second homology of $\mathcal{R}$ relative to the fence $\mathcal{F} \subset \mathcal{R}$ that defines $\partial\mathcal{D}$.

Theorem 3. *[1] For a set of nodes $\mathcal{X}$ in a domain $\mathcal{D} \subset \mathbb{R}^2$ satisfying Assumptions A1-A4, the sensor cover $\mathcal{U}$ contains $\mathcal{D}$ if there exists $[\alpha] \in H_2(\mathcal{R}, \mathcal{F})$ such that $\partial\alpha \neq 0$.*

The proof of this result is straightforward with an elementary knowledge of homology as in, say, Chapter 2 of [6]. We present an abbreviated proof.

Proof sketch. Define a simplicial realization map $\sigma : \mathcal{R} \to \mathbb{R}^2$ which sends vertices of $\mathcal{R}$ to the nodes $\mathcal{X} \subset \mathcal{D}$ and sends a k-simplex of $\mathcal{R}$ to the (potentially degenerate) k-simplex given by the convex hull of the vertices implicated. This σ takes the pair $(\mathcal{R}, \mathcal{F})$ to $(\mathbb{R}^2, \partial\mathcal{D})$. The long exact sequences on these two pairs yields the following commutative square:

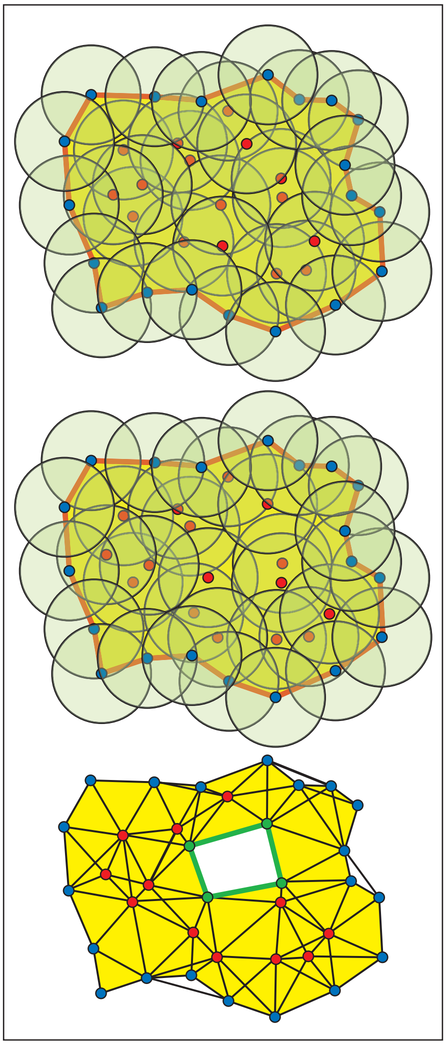

Figure 4. The homological criterion holds for some covers [top] but not for others [middle]. Failure is caused by a 1-cycle in the Rips complex [bottom].

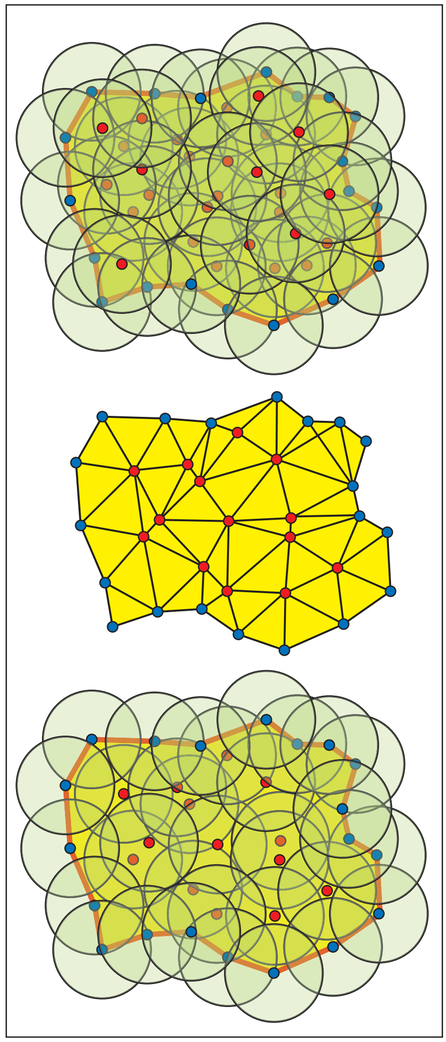

Figure 5. A redundant cover [top] can be simplified [bottom] by the appropriate choice of generator for $H_2(\mathcal{R},\mathcal{F})$ **[middle].**

$$
(1)\qquad
\begin{array}{ccc}
H_2(\mathcal{R},\mathcal{F}) & \xrightarrow{\delta_*} & H_1(\mathcal{F}). \\
\downarrow \sigma_* & & \downarrow \sigma_* \\
H_2(\mathbb{R}^2,\partial\mathcal{D}) & \xrightarrow{\delta_*} & H_1(\partial\mathcal{D})
\end{array}
$$

The homology class $\sigma_*\delta_*[\alpha]$ is the winding number of $\partial\alpha$ about $\partial\mathcal{D}$. Observe that $\sigma_*\delta_*[\alpha] = \sigma_*[\partial\alpha] \neq 0$, since, by assumption, $\partial\alpha \neq 0$. By commutativity of Equation (1), $\delta_*\sigma_*[\alpha] \neq 0$, and thus $\sigma_*[\alpha] \neq 0$.

Assume that $\mathcal{U}$ does not contain $\mathcal{D}$ and choose $p \in \mathcal{D} - \mathcal{U}$. Since, by the choice of r_c, every point in $\sigma(\mathcal{R})$ lies within $\mathcal{U}$, we have that $\sigma : (\mathcal{R},\mathcal{F}) \to (\mathbb{R}^2,\partial\mathcal{D})$ factors through the pair $(\mathbb{R}^2 - p, \partial\mathcal{D})$. However, $H_2(\mathbb{R}^2 - p, \partial\mathcal{D}) = 0$: contradiction via commutativity.

This homological criterion is sufficient but not necessary to verify coverage. The two networks illustrated in Figure 4 both cover the domain completely. Yet the homological criterion holds for one [top] and fails for the other [bottom]. The culprit in the case of failure is a cycle of length four in $H_1(\mathcal{R})$. This creates a hole in the Rips complex that is not present in the cover. Note, however, that a small change in the positions of the nodes implicated in this 4-cycle can create a hole in the cover without changing the topology of the network. No technique that relies solely upon the network topology can determine coverage in such a case. The homological criterion is effective for covers that are sufficiently robust with respect to perturbing the points while maintaining the network topology.

Generators for Power Conservation

The addition of some homological algebra to the sensor network can do more than confirm coverage. Indeed, it is a straightforward consequence of the proof that the domain $\mathcal{D}$ lies within the subcover of $\mathcal{U}$ given by those nodes implicated in the generator $[\alpha]$.

For a sensor network that has a highly redundant cover, one can save power and bandwidth by placing nonessential nodes in a *sleep mode*. The crucial question: which nodes can be deactivated without sacrificing coverage? Or, in a dynamic setting, how does one cycle nodes from sleep to wake modes without losing coverage? The answer lies in choosing the appropriate "minimal" generators for $H_2(\mathcal{R},\mathcal{F})$ that implicate as few 0-simplices as possible. Figure 5 gives an example of a "small" generator yielding a more efficient cover.

Pursuit and Evasion

There are a number of related contexts in which a homological criterion can solve a global problem. Consider the situation in which the nodes change position as a function of time. For simplicity, assume that the fence nodes are fixed. Such a situation might arise with sensors used to detect a forest fire, since one could establish a ring of fixed nodes outside the forest and allow the nodes inside the forest to be passively locomoted by environmental forces (e.g., animals).

It may well be the case that there are not enough sensors to cover the domain bounded by the outer ring. However, as the sensors change locations, holes in the cover can open and close in a complex fashion. The *evasion problem* for this scenario is whether an unknown *evader* can navigate through

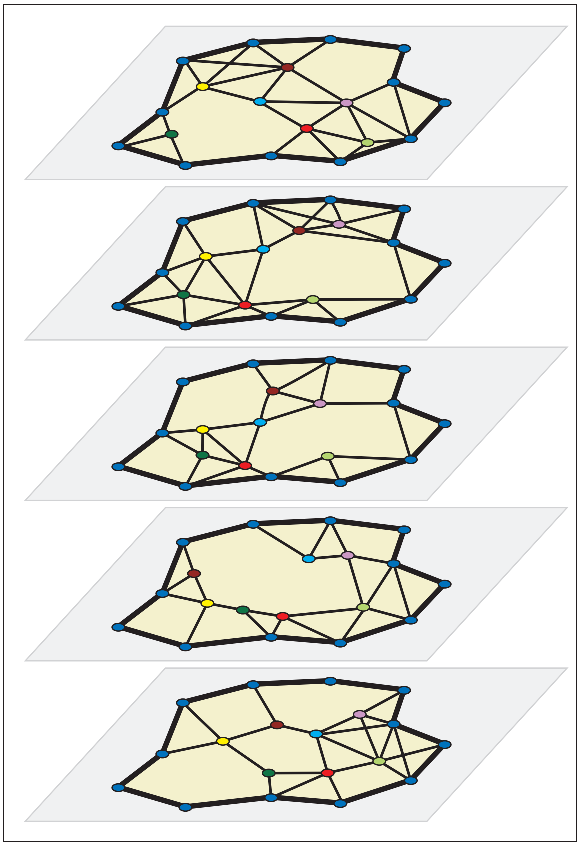

Figure 6. A time-sequence of network graphs for a mobile network. Does this network admit a wandering hole?

holes in the sensor cover without being detected. Even if coverage is never attained, one can still hope that any hole in which the evader begins is "squeezed" out eventually.

To address this problem, one proceeds as follows. Assume that the network communication graph is updated at certain time intervals $0 = t_1 < \ldots < t_i < \ldots < t_N = 1$, producing an ordered sequence of communication graphs Γ_i, for $i = 1 \ldots N$. These induce a corresponding sequence of Rips complexes $\mathcal{R}_i$. We impose the following additional assumptions:

A5: If two nodes are within broadcast radius at time steps t_i and t_{i+1}, then they remain so for all $t_i \le t \le t_{i+1}$.

A6: Nodes may go off-line or come on-line, represented by deleting or inserting the nodes in the appropriate graph Γ_i.

A7: Fence nodes remain fixed and on-line.

Given this sequence of network graphs (see Figure 6), it is by no means obvious whether there is a wandering hole in the coverage network. We amalgamate the sequence of Rips complexes into a single simplicial cell complex $\mathcal{AR}$ as follows. For each $i = 1, \ldots, N-1$, let $\mathcal{R}_i \cap \mathcal{R}_{i+1}$ denote the largest labeled subcomplex common to $\mathcal{R}_i$ and $\mathcal{R}_{i+1}$. This is well defined since all vertices (and thus all simplices) have unique labels. We define the *amalgamated Rips complex* to be the quotient of the disjoint union $\coprod \mathcal{R}_i$ obtained by identifying $\mathcal{R}_i \cap \mathcal{R}_{i+1} \subset \mathcal{R}_i$ with $\mathcal{R}_i \cap \mathcal{R}_{i+1} \subset \mathcal{R}_{i+1}$ for each i. This yields a cell complex built from simplices (though not necessarily a combinatorial simplicial complex, since multiple simplices may share the same vertex set). Note that, given A7, the fence $\mathcal{F}$ is a subcomplex of each $\mathcal{R}_i$ and thus is identified to a well defined cycle $\mathcal{F} \subset \mathcal{AR}$.

Theorem 4. *[1] Consider a set of mobile nodes $\mathcal{X}(t)$ in a domain $\mathcal{D} \subset \mathbb{R}^2$ satisfying A1-A7. Any continuous curve $p : [0,1] \to \mathcal{D}$ must have $p(t) \in \mathcal{U}(t)$ for some $0 \le t \le 1$ if there exists $[\alpha] \in H_2(\mathcal{AR}, \mathcal{F})$ such that $\partial\alpha \ne 0$.*

The proof of this result is in the same spirit as that of Theorem 3. Note that there are no bounds on the speed or cunning of the evader.

Persistence of Homology

The ease with which Theorem 3 is proved is due chiefly to the restrictions placed on the fence nodes in A4. With this condition, it is relatively easy to extend these results. Besides the time-dependent case reviewed above, homological methods work for domains that are not simply connected, for barrier coverage problems in 3-dimensions, for systems with communication errors or variable radii, and for hole detection and repair [1]. The control over the fence nodes is manifested in the proof of Theorem 3 in Equation (1), where $\sigma_* : H_1(\mathcal{F}) \to H_1(\partial\mathcal{D})$ is known to be an isomorphism.

Such control over the fence may be physically realistic in some settings where, say, one can explicitly build a ring of sensors around a potentially hazardous environment and then inject sensors in the interior of the domain. Equivalently, given an unbounded network and a cycle in the communication graph, one can query whether the region of the plane bounded by this cycle lies in the cover. A more realistic setting for boundary phenomena is one in which nodes can sense if they are near the boundary $\partial\mathcal{D}$ and can register themselves as fence nodes. For example, a very coarse range-finder can detect the presence of a wall within a set distance, without necessarily knowing the distance to the wall.

We therefore consider a system of stationary nodes which can detect the presence of the boundary of the domain $\partial\mathcal{D}$ within some fixed *fence radius* r_f. This choice of system leads to a collection of fence nodes $\mathcal{X}_f \subset \mathcal{X}$ which spans a *fence subcomplex* $\mathcal{F} \subset \mathcal{R}$, the maximal simplicial complex generated by the fence nodes and edges between them. The analogous coverage criterion in this case should be the existence of a generator $[\alpha] \in H_2(\mathcal{R}, \mathcal{F})$ such that $\partial\alpha \ne 0$. Unfortunately, this is no longer sufficient for coverage. Consider the network in Figure 7, in which the fence subcomplex $\mathcal{F}$ has a loop that is coned off to a disc in $\mathcal{R}$. This complex has $H_2(\mathcal{R}, \mathcal{F}) \neq 0$, yet the map $\sigma_* : H_1(\mathcal{F}) \to H_1(\partial\mathcal{D})$ is the zero-map, and Equation (1) is no longer useful in guaranteeing a cover. It is the existence of these *fake cycles* that complicates matters. To a "global" observer, the example of Figure 7 is easily seen to have degree zero. The challenge is to have the network determine this by "local" observations.

There is a simple homological criterion for coverage in this setting where the fence nodes are not controlled [2]: it uses *persistent homology* and requires some additional capabilities on the part of the sensor network. The heuristic behind this use of persistence is that the fake cycle of Figure 7 does not survive if the network increases its broadcast radius a small amount. Were this to happen, the "diagonals" of the 1-cycle in the fence subcomplex would be filled in, killing the relative 2-cycle.

We can generalize this one example to deal with arbitrary fake cycles by allowing for two broadcast radii: a "weak" and a "strong" signal. This also has the advantage of generalizing easily to compact domains $\mathcal{D} \subset \mathbb{R}^n$ for any $n \geq 2$. The precise assumptions are as follows:

P1: Nodes broadcast their unique ID numbers. Each node can detect the identity of any node within radius r_s via a *strong* signal, or via a *weak* signal within a larger radius r_w, where $r_w \geq r_s\sqrt{10}$.

P2: Nodes have radially symmetric covering domains of *cover* radius $r_c \geq r_s/\sqrt{2}$.

P3: Nodes lie in a compact domain $\mathcal{D} \subset \mathbb{R}^d$ and can detect the presence of the boundary $\partial\mathcal{D}$ within a *fence detection radius* r_f.

P4: The restricted domain $\mathcal{D} - C$ is connected, where

$$C = \left\{x \in \mathcal{D} : \|x - \partial\mathcal{D}\| \leq r_f + r_s/\sqrt{2}\right\}.$$

P5: The fence detection hypersurface $\{x \in \mathcal{D} : \|x - \partial\mathcal{D}\| = r_f\}$ has internal injectivity radius at least $r_s/\sqrt{2}$ and external injectivity radius at least r_s.

The crucial feature is that sensors that are within signal detection range can distinguish weak versus strong signals, yielding a binary measure of inrange distance. The fence nodes are not controlled, but there is a need for (somewhat severe) restrictions on the shape of the domain so as to exclude pinching (P4) and wrinkling (P5).

Such a system gives rise to a pair of Rips complexes, $\mathcal{R}_s$ and $\mathcal{R}_w$, computed at the strong and weak radii respectively. Each is outfitted with a fence subcomplex, $\mathcal{F}_s \subset \mathcal{R}_s$ and $\mathcal{F}_w \subset \mathcal{R}_w$. There is a natural inclusion of pairs

$$\iota : (\mathcal{R}_s, \mathcal{F}_s) \hookrightarrow (\mathcal{R}_w, \mathcal{F}_w), \tag{2}$$

since increasing the signal detection radius from r_s to r_w only increases network connectivity.

Theorem 5. *[2] For a set of nodes $\mathcal{X}$ in a domain $\mathcal{D} \subset \mathbb{R}^d$ satisfying P1-P5, the sensor cover $\mathcal{U}$ contains the restricted domain $\mathcal{D} - C$ if the induced homomorphism*

$$\iota_* : H_d(\mathcal{R}_s, \mathcal{F}_s) \to H_d(\mathcal{R}_w, \mathcal{F}_w)$$

is nonzero.

The key that makes this theorem work is a squeezing theorem for the Čech complex. For a set of points $\mathcal{X} \subset \mathbb{R}^d$, let $C_\epsilon(\mathcal{X})$ denote the Čech complex of the cover of $\mathcal{X}$ by balls of radius $\epsilon/2$. Let $\mathcal{R}_\epsilon(\mathcal{X})$ denote the Rips complex of the network graph having vertices $\mathcal{X}$ and edges between vertices within distance ϵ in $\mathbb{R}^d$.

Theorem 6. *[2] Fix $\mathcal{X}$ a set of points in $\mathbb{R}^d$. Given $\epsilon' < \epsilon$, There is chain of inclusions*

$$\mathcal{R}_{\epsilon'}(\mathcal{X}) \subset C_\epsilon(\mathcal{X}) \subset \mathcal{R}_\epsilon(\mathcal{X}) \quad \text{if} \quad \frac{\epsilon}{\epsilon'} \geq \sqrt{\frac{2d}{d+1}}.$$

Moreover, this ratio is the smallest for which the inclusions hold in general.

This is the type of result that is ideal for engineering applications. The Rips complex is computable, but does not give an accurate representation of the topology of the cover. The Čech complex gives the exact homotopy type of the cover, but it is not computable with the coarse information available from the network. Theorem 6 tells how to infer Čech data from Rips data.

This technique of comparison between Rips complexes at two different scales ϵ, ϵ' is a simple instance of the more general theory of persistent homology [3], [10]. This concerns the homological properties of nested families of topological spaces. Although the algebra and ideas involved are classical, the subject has been heavily driven by applications in computational geometry and nonlinear data analysis. Persistent homology is an algebraic topology for the twenty-first century.

Theorem 5 is not the final word in homological coverage criteria for systems with a fence radius and is best thought of as a proof-of-concept for homological methods. The hypotheses for this theorem flow from the mathematical details as opposed to the engineering details. For topological methods to make a serious contribution to security and sensor networks, it is important for the mathematics (and mathematicians) to work in conjunction with the engineers implementing the sensor networks.

The homological coverage criteria surveyed here are the beginning of a larger foray of topological ideas in the theories of networks and sensing. We note in particular the need for these coverage criteria to be distributed (so that networks can compute local homology and agree on global coverage), asynchronous (so that updates to the network are not dependent on a simultaneous sampling of the network), and fault tolerant (to accommodate the stochastic nature of sensor networks).

On Computational Topology

> "Topology! The stratosphere of human thought! In the twenty-fourth century it might possibly be of use to someone..."
>
> —*The First Circle,* A. Solzhenitsyn

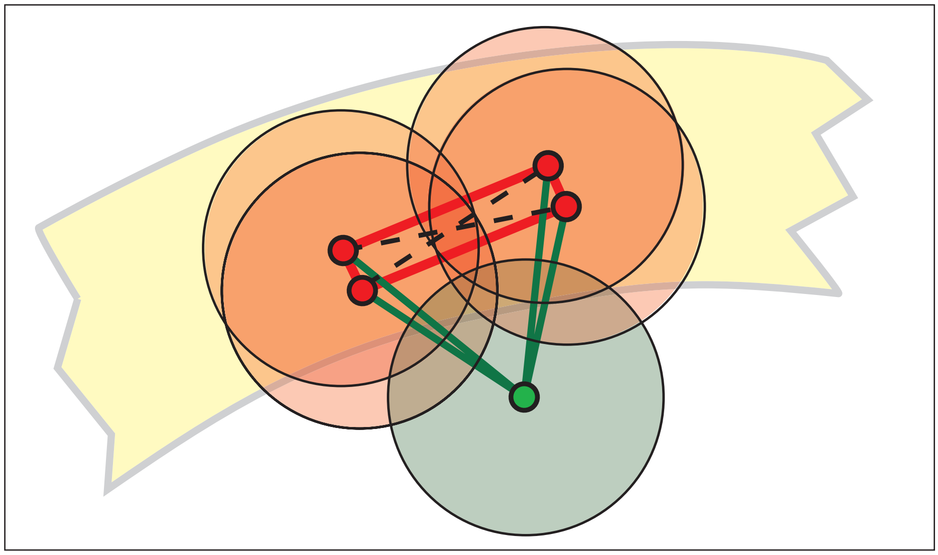

Figure 7. A fake relative 2-cycle in a system with a 1-cycle in the fence complex which is nullhomologous in the boundary collar.

The results we review here are but one branch of the rapidly evolving area of *applied computational topology*. The need to move from local to global is one that a large spectrum of engineers and scientists are finding to be prevalent. Very few of the calculus-based tools with which they are most familiar prove sufficient. Recently, it has been demonstrated that homology theory is useful for problems in data analysis and shape reconstruction, computer vision, robotics, rigorous dynamics from experimental data, and control theory. See [7] for an overview of some current applications.

Topology is especially keen at giving criteria for when one can or cannot find a particular global object (a homeomorphism, a nonzero section, an isotopy, etc.): this falls under the rubric of *obstruction theory*. This perspective is one that has not yet permeated the applied sciences, in which the question, *"What is possible?"* is usually approached from the top-down, *"Here's something we can build,"* as opposed to the bottom-up approach that topological methods yield. A brilliant example of this obstruction-theoretic viewpoint in an applied context is Farber's *topological complexity* for robot motion planning [5].

In this article, we use homology theory to give coverage criteria for networked sensors which are "nearly senseless". It seems counterintuitive that one can provide rigorous answers for a network with neither localization capabilities nor distance measurements. A topologist is not surprised that such coarse data can be integrated into a global picture. Some engineers are. Homological methods have the pleasant consequence that they may allow engineers to focus on designing simpler sensors that are nevertheless useful in a security network. Why bother miniaturizing GPS for "smart dust" if you can solve the problem without it? If topological methods can determine the minimal sensing needed to solve a global problem, then such methods may have significant impact on the way systems and sensors are developed and deployed.

References

[1] V. de Silva and R. Ghrist, Coordinate-free coverage in sensor networks with controlled boundaries via homology, to appear, *Intl. J. Robotics Research*.

[2] ____ , Coverage in sensor networks via persistent homology, conditionally accepted for publication, *Alg. & Geom. Top.*

[3] H. Edelsbrunner, D. Letscher, and A. Zomorodian, Topological persistence and simplification, *Discrete Comput. Geom.* **28** (2002), 511–33.

[4] D. Estrin, D. Culler, K. Pister, and G. Sukhatme, Connecting the physical world with pervasive networks, *IEEE Pervasive Computing*, **1**(1) (2002), 59–69.

[5] M. Farber, Topological complexity of motion planning *Discrete Comput. Geom.* **29**, no. 2, (2003), 211–21.

[6] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.

[7] T. Kaczynski, K. Mischaikow, and M. Mrozek, *Computational Homology,* Applied Mathematical Sciences 157, Springer-Verlag, 2004.

[8] H. Koskinen, On the coverage of a random sensor network in a bounded domain, *Proceedings of 16th ITC Specialist Seminar*, pp. 11-18, 2004.

[9] L. Vietoris, Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen, *Math. Ann.* **97** (1927), 454–72.

[10] A. Zomorodian and G. Carlsson, Computing Persistent Homology, *Discrete Comput. Geom.* **33** (2005), 249–74.