Generalized Coverage in Homological Sensor Networks*

Nicholas J. Cavanna,† Kirk Gardner‡
Donald R. Sheehy§

1 Introduction

In their seminal work on homological sensor networks, de Silva and Ghrist showed the surprising fact that it is possible to certify the coverage of a coordinate-free sensor network even with very minimal knowledge of the space to be covered [5]. We give a new, simpler proof of the de Silva-Ghrist *Topological Coverage Criterion (TCC)* that eliminates any assumptions about the smoothness of the boundary of the underlying space, allowing the results to be applied to much more general problems. The new proof factors the geometric, topological, and combinatorial aspects of this approach. This allows us to extend the TCC to support k-coverage, in which the domain is covered by k sensors, and weighted coverage, in which sensors have varying radii.

2 Background

Distances and Offsets. For a set X, let $\mathcal{P}(X)$ denote the power set of X and let $\binom{X}{k}$ denote the set of k-element subsets of X. Given a compact point set $A \subset \mathbb{R}^d$ with weights $w_a \geq 0$ for all $a \in A$ the **weighted distance** from a point x to a weighted point y is defined as the power distance

$$\rho_y(x)^2 := \|x-y\|^2 + w_y^2.$$

Such a set is referred to as a **weighted set**. We use weighted distances to model coverage by disks of varying radii, where larger weights correspond to smaller radii.

The **weighted k-nearest neighbor distance** from a point x to k points in a weighted compact set A is defined as

$$\mathrm{d}_k(x,A)^2 := \inf_{K \in \binom{A}{k}} \max_{y \in K} \rho_y(x).$$

Note that if $w_a = 0$ for all $a \in A$ and $k = 1$ then

$$\mathrm{d}_1(x,A) = \mathrm{d}(x,A) := \min_{y \in A} \|x-y\|.$$

Such as set is said to be **unweighted**.

The **canonical offsets** of a set A at a scale ε are defined as

$$A^\varepsilon := \{x \in X \mid \mathrm{d}(x,A) \leq \varepsilon\}.$$

We use the word “canonical” to distinguish these offsets from the **weighted (k,ε)-offsets** of a weighted compact set A, defined to be

$$A_k^\varepsilon := \{x \in X \mid d_k(x,A) \leq \varepsilon\}.$$

If A is unweighted we obtain the **(k,ε)-offsets**, the points within ε of k points in A. Note that for any weighted set A we have $A_k^\varepsilon \subseteq A^\varepsilon$. Thus, ε provides an upper bound on the radii.

Čech and Rips Complexes. The weighted **Čech complex** of a finite collection of points A in $\mathbb{R}^d$ at scale ε is defined as

$$\check{\mathrm{C}}\mathrm{ech}_\varepsilon(A) := \left\{\sigma \subseteq A \mid \exists x \in \mathbb{R}^d : \max_{p \in \sigma} \rho_p(x) \leq \varepsilon\right\}.$$

The **(Vietoris-)Rips complex** of A at scale ε is defined

$$\mathrm{Rips}_\varepsilon(A) := \left\{\sigma \subseteq A \mid \{p,q\} \in \check{\mathrm{C}}\mathrm{ech}_\varepsilon(A) \text{ for all } p,q \in \sigma\right\}.$$

The standard Rips and Čech complexes are obtained setting $w_p = 0$ for all $a \in A$.

An important result about the relationship of Čech and Rips complexes follows from Jung's Theorem [7] relating the diameter of a point set A and the radius of the minimum enclosing ball:

$$\check{\mathrm{C}}\mathrm{ech}_\varepsilon(A) \subseteq \mathrm{Rips}_\varepsilon(A) \subseteq \check{\mathrm{C}}\mathrm{ech}_{\vartheta_d \varepsilon}(A), \qquad (1)$$

where the constant $\vartheta_d = \sqrt{\frac{2d}{d+1}}$ for unweighted sets and $\vartheta_d = 2$ for weighted sets (see [1]).

The k-Barycentric Decomposition. Given a simplicial complex S we define a **flag** in S to be an ordered subset of simplices $\{\sigma_1, \ldots, \sigma_t\} \subset S$ such that $\sigma_1 \subset \ldots \subset \sigma_t$. The **barycentric decomposition** of S is the simplicial complex formed by the set of flags of S and is defined as $\mathrm{Bary}(S) := \{U \subset S \mid U \text{ is a flag of } S\}$. The vertices of the barycentric decomposition are the simplices of S. We define the **degree** of a flag $\sigma_1 \subset \cdots \subset \sigma_t$ to be $|\sigma_1|$.

Definition 1 *The **k-barycentric decomposition** of a complex S is defined*

$$k\text{-}\mathrm{Bary}(S) := \{U \subset S \mid U \textit{ is a flag in } S \textit{ of degree at least } k\}.$$

*This work was partially supported under grants CCF-1464379 and CCF-1525978.

†University of Connecticut
‡University of Connecticut
§University of Connecticut

The k-barycentric decomposition of the Čech and Rips complexes of a finite point set A at a scale ε are denoted $\check{\mathrm{C}}\mathrm{ech}_\varepsilon^k(A)$ and $\mathrm{Rips}_\varepsilon^k(A)$ respectively. By [8], we have the following results relating the k-barycentric decomposition of Čech and Rips complexes to the (k,ε)-offsets of their vertex set A.

Theorem 1 *Given a finite point set A, fixed k and any $\varepsilon \geq 0$ the k-barycentric decomposition of the Čech complex $\check{\mathrm{C}}\mathrm{ech}_\varepsilon^k(A)$ is homotopy equivalent to the (k,ε)-offsets A_k^ε.*

Theorem 2 *The k-barycentric decomposition of the Rips complex $\mathrm{Rips}_\varepsilon^k(A)$ is a ϑ_d-approximation to the (k,ε)-offsets A_k^ε.*

This result allows us to extend Equation 1 to the k-barycentric decomposition of the Čech and Rips complexes as follows:

$$\check{\mathrm{C}}\mathrm{ech}_\varepsilon^k(A) \subseteq \mathrm{Rips}_\varepsilon^k(A) \subseteq \check{\mathrm{C}}\mathrm{ech}_{\vartheta_d \varepsilon}^k(A), \qquad (2)$$

Homology and Persistent Homology. Homology is a tool from algebraic topology that gives a computable signature for a shape that is invariant under many kinds of topological equivalences. It gives a way to quantify the components, loops, and voids in a topological space. It is a favored tool for applications because its computation can be phrased as a matrix reduction problem with matrices representing a finite simplicial complex.

Throughout, we assume singular homology over a field, so the nth homology group $\mathrm{H}_n(C)$ of a space C is vector space. When considering the homology groups of all dimensions, we will write $\mathrm{H}_*(C)$. We will make extensive use of relative homology. That is, for a pair of spaces (A,B) with $B \subseteq A$, we write $\mathrm{H}_*(A,B)$ for the homology of A relative to B.

There are dual vector spaces to the homology groups called the **cohomology groups** and are denoted with superscripts as $\mathrm{H}^*(C)$. For finite-dimensional homology groups the **Alexander duality** [6] implies that for pairs of nonempty locally-contractible spaces in $\mathbb{R}^d \cup \{\infty\}$, the r-dimensional homology is isomorphic to the $(d-r)$-dimensional cohomology of the complement space, i.e.

$$\mathrm{H}_r(X,Y) \cong \mathrm{H}^{d-r}(\overline{Y},\overline{X}).$$

3 Geometric Assumptions

Strange examples abound in topology. One must make some assumptions about the underlying domain to make the TCC possible. In this section, we will introduce and illustrate the minimal geometric properties that we require of the bounded domain to be covered. Our goal is to weaken the geometric assumptions on the domain required for the topological coverage criterion of de Silva & Ghrist to apply it to domains without smooth manifold boundaries. In particular, for a domain $\mathcal{D}$ with boundary $\mathcal{B}$ satisfying the following conditions for $0 \leq 3\alpha \leq \beta$, we want to certify that a sample $P \subset \mathcal{D}$ covers $\mathcal{D}$ at scale α in the sense that $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subset P^\alpha$.

Assumptions

1. *(Non-degenerate)* $\mathcal{D}$ is compact, locally contractible, full dimensional in $\mathbb{R}^d$ and there exists $\varepsilon > 0$ such that $\mathcal{D} \hookrightarrow \mathcal{D}^\varepsilon$ induces a homotopy equivalence.
2. *(Components are not too small)* The map $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta} \hookrightarrow \mathcal{D} \setminus \mathcal{B}^{2\alpha})$ is surjective.
3. *(Components are not too close)* The map $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow D^{2\alpha})$ is injective.

Assumption 1 disallows degenerate cases in which some of the theorems listed in Section 2 cannot be applied. For example, the Alexander duality as we have stated it requires distinct, bounded, locally contractible spaces.

Assumptions 2 and 3 prevent components from appearing and disappearing in the inclusions $\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta} \hookrightarrow \mathcal{D} \setminus \mathcal{B}^{2\alpha}$ and $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow D^{2\alpha}$, respectively. These restrictions are in order to allow us to reliably compare the coverage region to the sampled subset of a disconnected domain in terms of the 0-dimensional homology, or connected components of related spaces. Specifically, Assumption 2 disallows domains with components that are too small to be included in the map from $\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta} \hookrightarrow \mathcal{D} \setminus \mathcal{B}^{2\alpha}$. Fig. 1 illustrates a domain in which the induced map is not surjective, allowing our algorithm to potentially report a false positive.

Assumption 3 requires that the components of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ are spaced out well enough so that no components are joined with inclusion into $\mathcal{D}^{2\alpha}$. This is required in order to be able to reliably bound the number of connected components of the shrunken domain by those of a computable combinatorial structure. Fig. 2 illustrates domains which violates Assumption 3 as components are lost in both $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow \mathcal{D}$ and $\mathcal{D} \hookrightarrow \mathcal{D}^{2\alpha}$.

Relationship to Geometric Assumptions in Prior Work

Previous work on the TCC by de Silva & Ghrist is restricted to connected domains with a smooth boundary in which a particular region around the domain has uniform thickness. This region is parameterized for smooth manifolds in the work by the injectivity radius, which serves to bound the region around the boundary in which no topological changes occur. The injectivity radius is closely related to the **reach** of a compact set K in $\mathbb{R}^d$, defined as the supremum of the distance r from K to any point x with a unique closest point $y \in K$

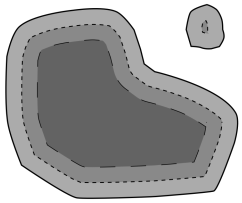

Figure 1: A domain that violates assumption 2 in which $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta} \hookrightarrow \mathcal{D} \setminus \mathcal{B}^{2\alpha})$ is not surjective as the upper-right component is pinched out in $\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta}$.

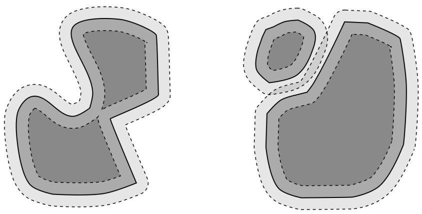

Figure 2: A domain that violates Assumption 3 in which $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow D^{2\alpha})$ is not injective as components are lost in both the inclusions $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow D$ and $\mathcal{D} \hookrightarrow D^{2\alpha}$.

such that $d(x, y) = d(x, K)$. For non-smooth compact sets K containing sharp corners, for example, note that the reach of K will be equal to zero, as for all $r > 0$ there must exist some $x \in \mathbb{R}^d \setminus K$ with at least two closest points in K a distance r from x approaching the sharp corner of the set.

This notion of feature size approaching a non-smooth feature is generalized in [2] as the **μ-reach**. Roughly speaking, the μ-reach parameterizes the reach in order to provide a meaningful measure of the region around a potentially non-smooth compact subset K in which no topological changes occur. In particular, the μ-reach is equal to the reach for $\mu = 1$, and converges to the minimum distance from K to the critical points of the distance function $d(\cdot, K)$ as μ approaches 0. This minimum distance is known as the **weak feature size** of $\mathcal{B}$ and was introduced in [3] as a way to parameterize compact sets that may not be smooth manifolds. For our purposes it can be understood as the minimum size of the topological features of a compact set. Thus, for a compact subset K of $\mathbb{R}^d$ we have that $\mathrm{reach}(K) \leq \mathrm{reach}_\mu(K) \leq \mathrm{wfs}(K)$.

We note that any bounded domain $\mathcal{D}$ such as that in Fig. 1 with a component in $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ in which the distance from every point in the component is at most $\alpha + \beta$ has a boundary with reach at most $\alpha + \beta$. Moreover, in the figures of 2 there exist points in $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ and $\mathcal{D}$ contained in distinct components within distance 2α of each other, respectively. It follows that the minimum distance to a critical point of $\mathcal{B}$, a point which lies in the convex hull of its nearest neighbors in $\mathcal{B}$, is at most 2α. Conversely, for any $\mathcal{D}$ such that the reach of $\mathcal{B}$ is strictly greater than $\alpha + \beta$ Assumption 2 is implied. Assumption 3 is implied for any bounded domain such that the weak feature size of its boundary is strictly greater than 2α. As $\mathrm{reach}(K) \leq \mathrm{wfs}(K)$ we therefore maintain our geometric assumptions for any domain $\mathcal{D}$ with a boundary $\mathcal{B}$ such that $\mathrm{wfs}(\mathcal{B}) > \alpha + \beta \geq 4\alpha$.

Relationship to the de Silva & Ghrist TCC For the sake of contrast, we will state the Topological Coverage Criterion as states by de Silva & Ghrist in [5]. Here we will assume points in a fixed point set P in a domain $\mathcal{D} \subset \mathbb{R}^d$ with boundary $\mathcal{B}$ have have uniform coverage radius r_c, fence detection radius (in which nodes can detect the boundary) r_f, and node-detection radii $r_s \leq \sqrt{2} r_c$, $r_w \geq r_s \sqrt{10}$. Moreover, we will let $Q = \{x \in P \mid d(x, \mathcal{B}) \leq r_f\}$ be the subset of P consisting of points sufficiently close to the boundary $\mathcal{B}$.

Theorem 3 (TCC (de Silva & Ghrist)) *Let P be a fixed set of nodes in a domain $\mathcal{D} \subset \mathbb{R}^d$ with boundary $\mathcal{B}$ such that each $p \in P$ the restricted domain $\mathcal{D} \setminus \mathcal{B}^{r_f + r_s/\sqrt{2}}$ is connected and the hypersurface $\Sigma = \{x \in \mathcal{D} \mid d(x, \mathcal{B}) = r_f\}$ has internal injectivity radius at least $r_s/\sqrt{2}$ and external injectivity radius at least r_s. The sensor cover P^{r_c} contains $\mathcal{D} \setminus \mathcal{B}^{r_f + r_s/\sqrt{2}}$ if the homomorphism*

$$\iota_* : H_d(\mathrm{Rips}_{r_s}(P), \mathrm{Rips}_{r_s}(Q)) \to H_d(\mathrm{Rips}_{r_w}(P), \mathrm{Rips}_{r_w}(Q))$$

induced by the inclusion $\iota : \mathrm{Rips}_{r_s}(P) \hookrightarrow \mathrm{Rips}_{r_w}(P)$ is nonzero.

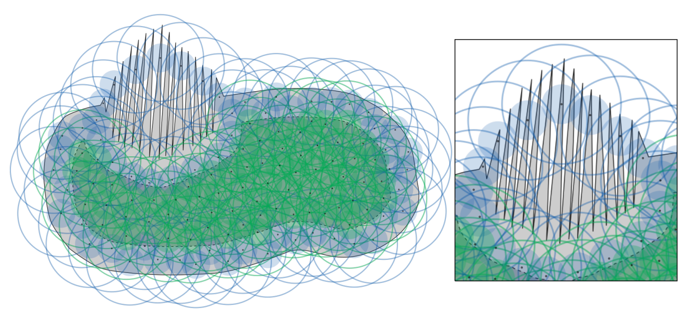

Figure 3: This instance illustrates the failure of Lemma 3.3 of [5] when the boundary is not smooth. A cycle that is trivial in the thickened boundary persists. This highlights the need to work with the relative homology of the domain modulo the boundary rather than the homology of the boundary alone. Such a cycle in the boundary cannot form a relative cycle.

According to [5, Remark 4.5], the smooth manifold hypothesis is a necessary requirement in order to apply the TCC without degenerating constants.[1] Because their analysis involves directly comparing this thickened region around boundary to this complex it was necessary to show that the thickness of this region is such that any topological noise in the complex is eliminated with inclusion from scale α to β. This amounts to proving that that cycles lying entirely in a thickening of the boundary of the domain cannot persist in the TCC, as in [5, Lemma 3.3]. Such a case is shown in Fig. 3, illustrating a domain without a smooth boundary in which the thickened boundary contains a cycle that persists across a range of scales. This is to be contrasted with our Lemma 5, requiring Assumptions 1 and 2, in which the persistence of this cycle does not indicate the persistence of a relative cycle. For example, although the domain in Fig. 3 clearly does not have a smooth boundary it does have weak feature size greater than $\alpha + \beta$, and therefore satisfies our geometric conditions. It is in this sense that Assumptions 2 and 3 serve to weaken the smoothness hypothesis in order to allow the TCC to be applied to domains, such as those with bounded weak feature size, which imply our minimal hypothesis.

4 The Generalized Topological Coverage Criterion

Consider a domain $\mathcal{D} \subset \mathbb{R}^d$, its boundary $\mathcal{B}$, and constants α and β such that $0 \leq 3\alpha \leq \beta$ satisfying Assumptions 1–3. Given a weighted finite point set $P \subseteq \mathcal{D}$ we will give a sufficient condition to guarantee the k-coverage of a shrunken domain $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$, where the k-covered region corresponds to the weighted offsets P_k^α as defined by d_k. We define $Q := P \cap \mathcal{B}^\alpha$, i.e. the subsample of P that is within distance α of the boundary.

Lemma 4 allows us to talk about the homology of the shrunken boundary in terms of relative homology where $\overline{A} := (\mathbb{R}^d \cup \infty) \setminus A$ denotes the **complement** of A in the compactification of $\mathbb{R}^d$ homeomorphic to the d-sphere S^d.

Lemma 4 *For all* $\varepsilon > 0$, $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^\varepsilon) \cong \mathrm{H}_0(\overline{\mathcal{B}^\varepsilon}, \overline{\mathcal{D}^\varepsilon})$.

Proof. Consider the inclusion $(\mathcal{D} \setminus \mathcal{B}^\varepsilon, \emptyset) \hookrightarrow (\overline{\mathcal{B}^\varepsilon}, \overline{\mathcal{D}^\varepsilon})$ and the corresponding map $\mathrm{H}_0((\mathcal{D} \setminus \mathcal{B}^\varepsilon, \emptyset) \hookrightarrow (\overline{\mathcal{B}^\varepsilon}, \overline{\mathcal{D}^\varepsilon}))$. For injectivity, given some non-trivial 0-chain $[x] \in \mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^\varepsilon)$, we can pick a representative point $x \in \mathcal{D} \setminus \mathcal{B}^\varepsilon \subseteq \overline{\mathcal{B}^\varepsilon}$. Given that $\mathcal{B}$ is the boundary of $\mathcal{D}$, a dimension-n space, then there exists no paths from $\mathcal{D} \setminus \mathcal{B}^\varepsilon$ to $\overline{\mathcal{B}}$, so $[x] \neq 0 \in \mathrm{H}_0(\overline{\mathcal{B}^\varepsilon}, \overline{\mathcal{D}^\varepsilon})$. For surjectivity, given some $[x] \in \mathrm{H}_0(\overline{\mathcal{B}^\varepsilon}, \overline{\mathcal{D}^\varepsilon})$, it represents a point on a connected component on $\overline{\mathcal{B}^\varepsilon} \setminus \overline{\mathcal{D}^\varepsilon} = \mathcal{D} \setminus \mathcal{B}^\varepsilon$, and thus a homology class $[x] \in \mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^\varepsilon)$. □

[1] [5, Remark 4.5] states that, although domains with a polygonal boundary are admissible in practice, the constant r_w would blow up along with the angle of the sharpest corner of the outermost boundary component.

We will assume non-negative weights $w_x \geq 0$ assigned to each $x \in P$, and that $w_x = 0$ for all points $x \in \mathcal{D} \setminus P$. This implies that $\mathcal{D}_k^\varepsilon = \mathcal{D}^\varepsilon$, and similarly $\mathcal{B}_k^\varepsilon = \mathcal{D}^\varepsilon$, so we will simply use the notation $\mathcal{D}^\varepsilon$ and $\mathcal{B}^\varepsilon$ throughout. Moreover, we know that $P_k^\alpha \subseteq \mathcal{D}_k^\alpha = \mathcal{D}^\alpha$ by the monotonicity of d_k. For any arbitrary weighted compact set $A \subseteq \mathcal{D}$, $A_k^\varepsilon \subseteq A_1^\varepsilon \subseteq A^\varepsilon$ and $Q \subseteq B^\alpha$, for $\varepsilon \geq 0$, $Q_k^\varepsilon \subseteq Q^\varepsilon \subseteq B^{\alpha+\varepsilon}$.

Diagram (3) relates the connected components of the pairs, and induces the map $\pi_* : \mathrm{im}\ j_* \to \mathrm{im}\ i_*$.

$$\begin{array}{ccc} \mathrm{H}_0(\overline{\mathcal{B}^{\alpha+\beta}}, \overline{\mathcal{D}^{\alpha+\beta}}) & \xrightarrow{j_*} & \mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}}) \\ \downarrow & & \downarrow \\ \mathrm{H}_0(\overline{Q_k^\beta}, \overline{P_k^\beta}) & \xrightarrow{i_*} & \mathrm{H}_0(\overline{Q_k^\alpha}, \overline{P_k^\alpha}) \end{array} \tag{3}$$

Though reversed and inverted by the dualities, this map describes the topology of the offsets embedded into the domain, where the scale change eliminates noise. That is, it captures exactly the topological information we want. Analyzing π_* directly simplifies the proof and aids in eliminating some hypotheses.

The following two lemmas prove two important properties of π_*. These will be used to give a computable way to infer coverage from the rank of i_*.

Lemma 5 *Given Assumptions 1 and 2, the map* π_* *is surjective.*

Proof. Assumption 2 implies that j_* is surjective by Alexander Duality. We choose a basis for im i_* such that each basis element is a point in $P_k^\alpha \setminus Q_k^\beta$. Consider $x \in P_k^\alpha \setminus Q_k^\beta$ such that $[x] \neq 0 \in \mathrm{im}\ i_*$. If $x \in \overline{\mathcal{B}^{2\alpha}}$, then $x \in \mathcal{D} \setminus \mathcal{B}^{2\alpha}$ so $[x] \neq 0 \in \mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}})$. Because j_* is surjective, $\mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}}) = \mathrm{im}\ j_*$ and thus $\pi_*([x]) = [x]$ and so $[x] \in \mathrm{im}\ \pi_*$.

If $x \in \mathcal{B}^{2\alpha}$, then there is a point $y \in \mathcal{B}$ such that $\|x - y\| \leq 2\alpha$. Because $x \in \overline{Q_k^\beta}$ by hypothesis, $\mathrm{d}_k(x, Q) > \beta$. We will show that if a point z is in the line segment $\overline{xy}$, then $z \in \overline{Q_k^\alpha}$. For any $z \in \overline{xy}$, we have $\|x - z\| \leq \|x - y\| \leq 2\alpha$. So,

$$\begin{aligned} \mathrm{d}_k(z, Q) &\geq \mathrm{d}_k(x, Q) - \|x - z\| && [\mathrm{d}_k \text{ is Lipschitz}] \\ &> \beta - 2\alpha && [\mathrm{d}_k(x, Q) > \beta \text{ and } \|x - z\| \leq 2\alpha] \\ &\geq \alpha && [\beta \geq 3\alpha] \end{aligned}$$

So, we conclude that $z \in \overline{Q_k^\alpha}$, and thus $\overline{xy} \subseteq \overline{Q_k^\alpha}$.

The definition of Q implies that $\mathcal{B} \cap \overline{Q_k^\alpha} \subseteq \overline{P_k^\alpha}$, and so $y \in \overline{P_k^\alpha}$. Any path $\gamma : [0, 1] \to \overline{Q_k^\alpha}$ such that $\gamma(0) = x$ and $\gamma(1) = y$, generates a class $[\gamma]$ in the chain group $C_1(\overline{Q_k^\alpha})$ containing γ. For $[\gamma] \in C_1(\overline{Q^\alpha}_k, \overline{P_k^\alpha})$ it follows $\partial([\gamma]) = [x + y] = [x]$ as $y \in \overline{P_k^\alpha}$, and therefore that there must exist $z \in \overline{xy} \cap Q_k^\alpha$. This is a contradiction as we have shown that $\overline{xy} \cap Q_k^\alpha = \emptyset$, and thus we conclude x cannot be in $\mathcal{B}^{2\alpha}$. □

The following lemma therefore allows us to confirm coverage by comparing the ranks of im i_* and im j_*.

Lemma 6 *Given Assumption 1, if π_* is injective then $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$.*

Proof. The proof is essentially the same as that presented by de Silva & Ghrist [5]. We include it here in our own notation for completeness.

We will prove this by contradiction. Assume there exists $x \in (\mathcal{D} \setminus \mathcal{B}^{2\alpha}) \setminus P_k^\alpha$,and thus $[x] \neq 0 \in \mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}})$. This is true because as we know that x is in the interior of $\mathcal{D}$, so it is on a connected component of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$. Consider the following sequence:

$$\mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}}) \xrightarrow{f_*} \mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}} \cup \{x\}) \xrightarrow{g_*} \mathrm{H}_0(\overline{Q_k^\alpha}, \overline{P_k^\alpha})$$

As $f_*([x]) = 0 \in \mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}} \cup \{x\})$, then $(g_* \circ f_*)([x]) = 0$. But we have a contradiction as $g_* \circ f_* = \pi_*$, and $\pi_*([x]) \neq 0$ by injectivity, so $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$. □

As Lemma 5 asserts that π_* is surjective under our assumptions, Lemma 6 can therefore be used to confirm coverage by providing conditions in which π_* is injective. Thus, the following theorem provides sufficient conditions to confirm $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$. Note that it will not yet give us an algorithm (that will come in Theorem 10), but instead gives a result about the offsets directly rather than an embedding of a Rips complex as was used in previous work.

Theorem 7 (Geometric TCC) *Consider $\mathcal{D} \subset \mathbb{R}^d$ with boundary $\mathcal{B}$ satisfying Assumptions 1 and 2. Let α and β be constants such that $0 < 3\alpha \leq \beta$. Let $P \subset \mathcal{D}$ be a finite set with $Q = P \cap \mathcal{B}^\alpha$. Let i_* and j_* be the maps in Diagram* (3)*. If* rk $i_* \geq$ rk j_* *then $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$.*

Proof. Given Assumptions 1 and 2, Lemma 5 implies that $\pi_* :$ im $j_* \to$ im i_* is surjective, and so rk $i_* \leq$ rk j_*. By hypothesis, rk $i_* \geq$ rk j_*, so it follows that rk $i_* =$ rk j_*. Because both the images are finite-dimensional, π_* is an isomorphism, and therefore it is injective. Lemma 6 then implies $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$. □

5 Computing the TCC

In the previous section we prove sufficient conditions for generalized coverage in terms of the offsets of the input points. However, we may not be able to compute these offsets, because we do not know the positions of the points in P. Instead, we use Rips complexes in the algorithm.

Let $\mathrm{Rips}_\beta(X)$ and $\check{\mathrm{C}}\mathrm{ech}_\beta(X)$ denote respectively the Rips and Čech complexes of a set X at scale β. Let R_β^k be the pair of k-barycentric Rips complexes $(\mathrm{Rips}_\beta^k(P), \mathrm{Rips}_\beta^k(Q))$ and let C_β^k be the pair of k-barycentric Čech complexes $(\check{\mathrm{C}}\mathrm{ech}_\beta^k(P), \check{\mathrm{C}}\mathrm{ech}_\beta^k(Q))$ as defined in Section 2. If $k = 1$ and P is unweighted we define the standard Rips and Čech complex pairs $R_\beta^1 := (\mathrm{Rips}_\beta(P), \mathrm{Rips}_\beta(Q))$ and $C_\beta^1 := (\check{\mathrm{C}}\mathrm{ech}_\beta(P), \check{\mathrm{C}}\mathrm{ech}_\beta(Q))$.

Algorithm 1 is for checking k-coverage of the shrunken domain by a weighted sample P, i.e. that $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$. The algorithm requires that the point samples each of the connected components of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$. It first constructs three Rips complexes based on the input parameters (α, β, P, Q, k): $\mathrm{Rips}_\alpha(P)$, R_{α/ϑ_d}^k and R_β^k. It then checks a condition relating the homology of the complexes, and if it satisfied, k-coverage is guaranteed. Note that if the algorithm's output is false it does not necessarily mean there is not coverage. Lemma 8, Lemma 9 and Theorem 10 together provide a proof of correctness of Algorithm 1. Lemma 8 bounds the rank of the map between the Rips complexes at different scales by rk i_*, in order to compare it to rank j_* through Theorem 7. Lemma 9 states that if the components are separated enough, formally defined in Assumption 3, then the number of connected components of the Rips complex at scale α of P provides an upper bound for the number of components of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$.

Algorithm 1 Check if $\mathcal{D} \setminus B^{2\alpha} \subseteq P_k^\alpha$

procedure K-COVERAGE(α, β, P, Q, k)
    construct $\mathrm{Rips}_\alpha(P)$
    let $c :=$ dim $H_0(\mathrm{Rips}_\alpha(P))$
    construct R_{α/ϑ_d}^k and R_β^k
    let $r :=$ rk $\mathrm{H}_d(R_{\alpha/\vartheta_d}^k \hookrightarrow R_\beta^k)$
    if $c = r$ **then return** True
    else return False

Lemma 8 *The rank of the map $\mathrm{H}_d(R_{\alpha/\vartheta_d}^k \hookrightarrow R_\beta^k)$ induced by inclusion is at most* rk i_*.

Proof. For the case of $k = 1$, the Persistent Nerve Lemma [4] says that for $\varepsilon \geq 0$, $\mathrm{H}_*(C_\varepsilon^1) \cong \mathrm{H}_*(P_1^\varepsilon, Q_1^\varepsilon)$. The Universal Coefficient Theorem with respect to Diagram (3) implies that $\mathrm{rk}(\mathrm{H}_d(C_\alpha^1 \hookrightarrow C_\beta^1)) = \mathrm{rk}\ i_*$. Moreover, the inclusion $R_{\alpha/\vartheta_d}^1 \hookrightarrow R_\beta^1$ can be factored as

$$R_{\alpha/\vartheta_d}^1 \hookrightarrow C_\alpha^1 \hookrightarrow C_\beta^1 \hookrightarrow R_\beta^1.$$

It follows that

$$\mathrm{rk}(\mathrm{H}_d(R_{\alpha/\vartheta_d}^1) \to \mathrm{H}_d(R_\beta^1)) \leq \mathrm{rk}(\mathrm{H}_d(C_\alpha^1) \to \mathrm{H}_d(C_\beta^1)) = \mathrm{rk}\ i_*.$$

For $k \geq 2$, Theorem 2 states that $(\mathrm{Rips}_\varepsilon^k(P), \mathrm{Rips}_\varepsilon^k(Q))$ is a ϑ_d-approximation to $(P_k^\varepsilon, Q_k^\varepsilon)$. This implies that $\mathrm{H}_*(R_{\varepsilon/\vartheta_d}^k) \cong \mathrm{H}_*(P_k^\varepsilon, Q_k^\varepsilon)$, so the previous argument naturally follows for these cases as well. □

Lemma 9 *Given P has at least one point on each connected component of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$, if Assumptions 1 and 3 are satisfied then the number of connected components of $\mathrm{Rips}_\alpha(P)$ is greater than or equal to the number of connected components of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$.*

Proof. Assume there exists $p, q \in P$ such that p and q are connected in $\mathrm{Rips}_\alpha(P)$, but not in $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$. This implies that $\|p - q\| \leq 2\alpha$ and $[p] \neq [q]$ in $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{2\alpha})$. However, $\overline{pq} \in \mathcal{D}^{2\alpha}$ as the distance between p and q is less than 2α, so $[p] = [q]$ in $\mathrm{H}_0(\mathcal{D}^{2\alpha})$, which implies that $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow \mathcal{D}^{2\alpha})$ is not injective, a contradiction to Assumption 3. □

Theorem 10 (Algorithmic TCC) *Consider a domain $\mathcal{D} \subset \mathbb{R}^d$ with boundary $\mathcal{B}$ and constants α, β, where $0 \leq 3\alpha \leq \beta$, satisfying Assumptions 1, 2 and 3. Let $P \subset \mathcal{D}$ be a finite point sample, $|P| \geq \max\{k, m\}$, where $m = \mathrm{H}_0(D \setminus \mathcal{B}^{2\alpha})$, such that there is a point $p \in P$ in each of the m connected components of $D \setminus \mathcal{B}^{2\alpha}$. If $\mathrm{rk}\ \mathrm{H}_d(R^k_{\alpha/\vartheta_d} \hookrightarrow R^k_\beta) = \dim\ H_0(\mathrm{Rips}_\alpha(P))$, then $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P^\alpha_k$.*

Proof. For simplicity, define $a_* := \mathrm{H}_d(R^k_{\alpha/\vartheta_d} \hookrightarrow R^k_\beta)$ and set $c = \dim\ \mathrm{H}_0(\mathrm{Rips}_\alpha(P))$. By our hypothesis and Lemma 8, $\mathrm{rk}\ i_* \geq \mathrm{rk}\ a_* = c$. By Lemma 9, $c \geq m$, and Assumption 2 implies that j_* is surjective by Alexander duality, so $m = \mathrm{rk}\ j_*$. Thus $\mathrm{rk}\ i_* \geq \mathrm{rk}\ a_* = c \geq m = \mathrm{rk}\ j_*$, namely $\mathrm{rk}\ i_* \geq \mathrm{rk}\ j_*$, so by Theorem 7 we can conclude $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P^\alpha$. □

From this algorithm we can see that, even if we do not know the number of connected components of D_0, as long as we know which components have been sampled we can provide a condition to certify coverage of the subdomain that P has been sampled from.

6 Conclusion

The TCC gives an effective algorithm for certifying coverage of coordinate-free sensors in an unknown domain. In this paper, we generalized the TCC to certify coverage in spaces whose boundaries may not be smooth. We replaced the smoothness assumption with much weaker conditions, that the domain is non-degenerate in some sense (Assumption 1), that the components are not too small (Assumption 2), and that the components are not to close (Assumption 3).

Although the language of homological sensor networks might imply that the application is restricted to sensors, we hope that the more general geometric conditions provided in this paper will lead to applications in data analysis. Specifically, eliminating the smoothness assumption should make this approach amenable to real data problems.

References

[1] Mickaël Buchet, Frédéric Chazal, Steve Y. Oudot, and Donald R. Sheehy. Efficient and robust persistent homology for measures. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 168–180, 2015.

[2] Frédéric Chazal, David Cohen-Steiner, and André Lieutier. A sampling theory for compact sets in euclidean space. *Discrete & Computational Geometry*, 41(3):461–479, 2009.

[3] Frédéric Chazal and André Lieutier. Weak feature size and persistent homology: Computing homology of solids in R^n from noisy data samples. In *Proceedings of the 21st ACM Symposium on Computational Geometry*, pages 255–262, 2005.

[4] Frédéric Chazal and Steve Yann Oudot. Towards persistence-based reconstruction in euclidean spaces. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, SCG '08, pages 232–241, New York, NY, USA, 2008. ACM.

[5] Vin de Silva and Robert Ghrist. Coverage in sensor networks via persistent homology. *Algebraic & Geometric Topology*, 7:339–358, 2007.

[6] William Julian, Ray Mines, and Alexander. Alexander duality. 106:115–127, 1983.

[7] Heinrich Jung. Über die kleinste kugel, die eine räumliche figur einschließt. *J. Reine Angew. Math*, 123:214–257, 1901.

[8] Donald R. Sheehy. A multicover nerve for geometric inference. In *The Canadian Conference in Computational Geometry*, pages 309–314, 2012.