Localized Homology*

Afra Zomorodian[†] Gunnar Carlsson[‡]

Preliminary Draft — 12/2/05

Abstract

In this paper, we introduce *localized homology*, a theory for finding local geometric descriptions for topological attributes. Given a space and a cover of subspaces, we construct the blowup complex, a derived space that contains both local and global information. The persistent homology of the blowup complex localizes the topological attributes of the space. Our theory is general and applies in all dimensions. After an informal description, we formalize our approach for general spaces, adapt it for simplicial complexes, and develop a simple algorithm that works directly on the input. In each stage, we prove the theoretical equivalence of the methods. We also implement our algorithm and give preliminary results to validate our methods in practice.

1 Introduction

In this paper, we address the problem of localizing topological attributes. Simply put, the question is whether we can distinguish between the solid and dashed loops in Figure 1. They both describe the same tunnel, but one is *ugly!*

Topology describes how a space is connected. This information reflects the presence of certain *qualitative* features in the space, such as the tunnel in the figure. While algebraic topology, and specifically the formalism of *homology*, is capable of detecting the *existence* of such features, it cannot directly tell us about their *location*. For example, homology may produce the long solid loop as a description of the tunnel in the figure. While topologically correct, this description is geometrically useless. The *localization problem* is determining the location of topological features within a space.

Algebraic topology is about functoriality. The critical insight here is that the recent theory of persistent homology [9] is a computational view of functoriality. Our main contribution in this paper is the synthesis of the classic Mayer-Vietoris blowup construction [6] with persistent homology, giving us a solution to localization. Unlike prior work, our theory applies in all dimensions: we may use the same method to localize tunnels, surfaces that enclose empty spaces, or arbitrary-dimensional features in arbitrary-dimensional spaces. We derive a simple practical algorithm, implement it, and show results at the end of this paper to substantiate our approach.

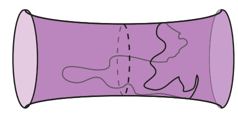

Figure 1. The localization problem. Both loops describe the tunnel. But homology cannot distinguish between them.

The focus of this paper is on the theoretical foundations of our work. In this section, we begin by briefly describing what topological attributes are, and why we are interested in both discovering and localizing them. We next familiarize the reader with the difficulties associated with localizing homology. To localize, we integrate geometry into homology to arrive at our theory. We end this introduction with a complete non-theoretical presentation of the key ideas in this paper using a simple example. In the sections that follow, we formalize the intuition.

1.1 The Problem

Most disciplines examine the *geometry* of a space, focusing on quantitative questions and local properties. The topology of a space, however, may have significant repercussions on the ability of geometric algorithms to perform effectively or even terminate. In *computer graphics*, undersampling and noise often result in extraneous topology, such as the spurious handles in the Stanford Buddha surface in Figure 2 on the left. This false connectivity hinders subsequent geometry processing, such as simplifica-

*Research partially supported by DARPA under grant 32905.

[†]Department of Computer Science, Stanford University, Stanford, California.

[‡]Department of Mathematics, Stanford University, Stanford, California.

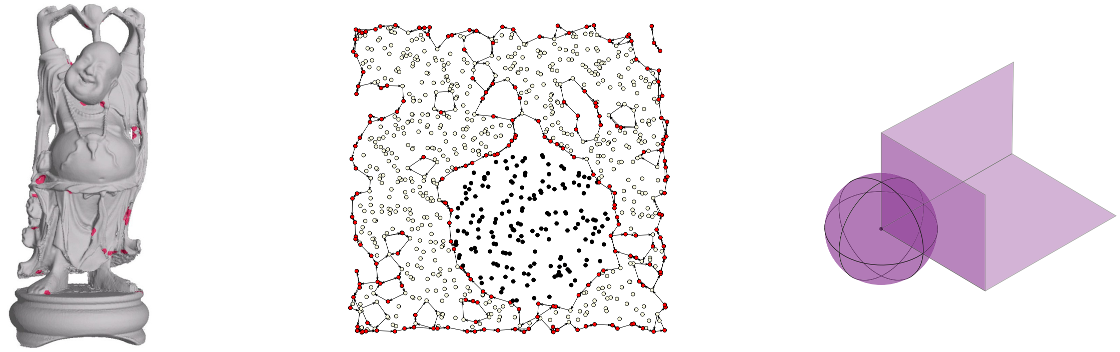

Figure 2. The localization problem emerges. Left: The Buddha isosurface has genus 104, instead of the expected 6. The insidious handles exist in the highlighted regions and hamper subsequent geometry processing [7]. Center: The failed black sensors in the sensor network result in a large hole in network that breaks the geographical greedy forwarding method for routing [3]. Right: We may identify the corner point through the topology of its fiber: three intersecting circles, instead of one for a smooth point [2].

tion, smoothing, and parametrization. In *sensor networks*, nonuniform distribution, terrain features, or catastrophic failure of nodes may lead to regions without working sensors, as shown in Figure 2 in the center. These holes break efficient but greedy communication algorithms [3]. In *robotics*, we require a compact representation of the *configuration space* of a robot for the fast computation of ensemble properties, such as the probability of folding *p-fold* of a protein conformation [1]. In *shape description*, we need geometric descriptions of topological attributes of the tangent complex in order to identify singular features, such as corner points or edges, as shown in Figure 2 on the right [2].

In each case, the existence of the topological attributes like holes and handles cause complications. And the resolution of these complications require localized descriptions. The emergence of topological questions in many areas has given rise to the area of *computational topology*, the area to which this question belongs [8].

Homology is the algebraic invariant often used to capture topological attributes since it is easily computable in all dimensions. This method characterizes the topology of a space through the structure of its holes [5]. It extends the notion of a hole or *cycle* to all dimensions. A cycle has an intuitive meaning in $\mathbb{R}^3$: A 0-*cycle* is a component (piece) of the space. A 1-*cycle* is a loop that goes around a tunnel. And a 2-*cycle* is a surface that encloses an empty space. A homology cycle is really a class of equivalent *homologous* cycles, all of which characterize the same topological attribute. In each dimension k, the cycles interact to form a vector space of cycles H_k. Any basis for this vector space H_k has the same rank, the *Betti number* β_k of the space.

While effective in capturing topology, homology is inherently nonlocal. We already had a glimpse of this nonlocality in Figure 1. The nonlocality, however, is more fundamental. A homology cycle may have multiple components like the 1-cycle in Figure 3(a). Moreover, homology computes a basis without regard to geometry, so any cycle in the 1-skeleton of a tetrahedron in Figure 3(b) is a candidate basis element and not those that are geometrically local. Similarly, for the graph in Figure 4, homology may choose one of the nonlocal bases instead of the local one. The examples demonstrate that homology does not favor localization by nature. It has no knowledge of the geometry of the space and cannot identify local bases.

1.2 Adding Geometry

We need to incorporate geometrical selection into homology. An immediate idea is to compute the topology of local pieces of the space. Suppose we *cover* the graph in Figure 4 with a number of local sets whose union contains the graph, as shown in Figure 5(a). Computing the homology within the sets gives us the desired local basis in Figure 4: Each set contains one cycle and putting the two cycles together yields the local basis.

We must be careful in assembling the cover, however,

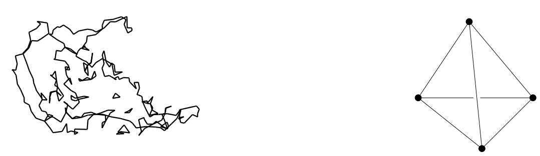

Figure 3. Nonlocality. Left: A homology cycle may have multiple components. Right: A tetrahedron has four faces, but its 1-skeleton has cycles that form a vector space of rank $\beta_1 = 3$.

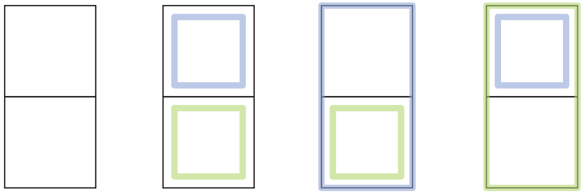

Figure 4. Bases. A graph (left) and three possible bases for its H_1. Only the first basis to its right is local.

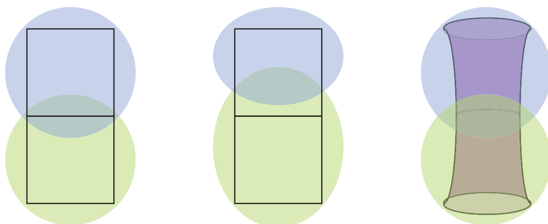

Figure 5. Covers. (a) In the good cover, each cycle is localized in a single set, resulting in the local basis in Figure 4. (b) In the bad cover, the top cycle is in neither set, so local computation fails to see it. (c) In the ugly cover, the tunnel appears in both sets and is discovered twice.

as it may overlook topological attributes. For instance, the cover in Figure 5(b) localizes only the bottom cycle. Neither set in the cover contains the cycle on the top, so the local computation fails to see it. This failure is benign, however, as it is a consequent of the cover. We may detect it easily by comparing the topology of the entire space to the result of the local computation and patching our cover.

A more distressing problem emerges when a topological attribute appears in multiple sets in the cover, as the tunnel in Figure 5(c). Locally, the tunnel is discovered twice, once in each set. Globally, the discovered 1-cycles are equivalent as they both go around the same tunnel. To recover the equivalence of the two tunnels, we need to understand how the local pieces are glued to each other within the intersection of the two cover sets. The theoretical gadget for exposing this relationship is the *Mayer-Vietoris sequence* [5]. For any cover, the sequence relates the homology of a space to the homology of the pieces within the cover sets and their intersections. Unfortunately, the sequence is only useful for computation by hand. It does not yield an implementable algorithm for arbitrary spaces.

1.3 Our Approach

Our approach may be viewed as a computational version of the Mayer-Vietoris sequence. We begin by applying the idea from the previous section, blowing up the space into local pieces according to the cover. For example, the graph containing three cycles covered by two sets in Figure 6(a) is blown up into two pieces in (b), each with two 1-cycles. Since the middle cycle of the original space is contained in the intersection of the cover sets, it exists in both local pieces. To recover the global topology, we equate the two copies of the middle cycle by gluing a cylinder to them. The resulting *Mayer-Vietoris blowup complex* in Figure 6(c) has the same number of cycles as the original space but also incorporates the geometric cover information within its structure.

We now need to compute homology bases for the blowup complex that are compatible with bases for the local pieces. Fortunately, the theory of *persistent homology* furnishes the required bases [9]. We incrementally assemble the blowup complex so that the local pieces are included at time 0 and the cylinder is sewn in at time 1, completing the structure. Persistence computes compatible homology bases across this growth history. Therefore, it can track individual basis elements, representing their lifetimes in a multiset of intervals called a barcode. The barcode for our example, shown in Figure 6(d), has three half-infinite intervals, corresponding to the three 1-cycles in both the original space and its blowup complex. But we can also color the barcode to show where the 1-cycles are located. There are four intervals at time ≤ 1, representing the four local 1-cycles in Figure 6(b). At time 1, the cylinder equates the two copies of the middle 1-cycle, so one of the two intervals that represent the two copies ends. The choice of the interval corresponds to the choice of the basis representative of the middle cycle lying in either of the two sets of the cover. As the two are homologous, the choice is arbitrary.

To summarize, given a space equipped with a cover, we incorporate the geometry contained within the cover into homology by building the blowup complex and computing its persistent homology. We call this method *localized homology*.

1.4 Outline

We believe an important aspect of our approach is the clear separation of geometry and topology. In this paper, we focus purely on topology: the theory of localized homology and an algorithm for its computation. We emphasize that different problems will require different covers and a variety of algorithms may be employed to select

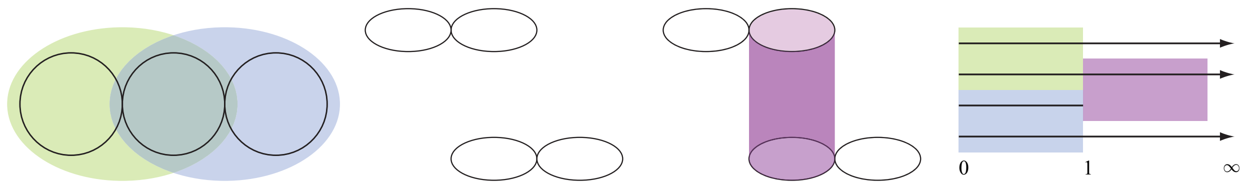

(a) Space and cover (b) Local pieces ($t = 0$) (c) Blowup ($t = 1$) (d) Persistence barcode

Figure 6. Our approach. Given a space equipped with a cover (a), we first blow up the space into local pieces (b) and then glue back the pieces to get the blowup complex (c), giving us a filtration of two complexes at times $t = 0$ and $t = 1$. The persistent homology of the blowup complex gives us a barcode (d) that localizes the topology of the original space with respect to the given cover.

appropriate covers. Moreover, a hierarchical approach toward localization is feasible: if we can localize a cycle within a set in any cover, we can improve the cycle description by recursively localizing within that set. But regardless of the origin of the cover, we need to compute localized homology. This is our task in the rest of this paper. In Section 2, we present the algebraic concepts utilized in the paper. Section 3 contains the contributions of our paper. We begin with a formal definition of localized homology for general spaces in Definition 5. We then provide alternate definitions for combinatorial spaces. These definitions are computationally feasible and we prove that they provide the same answer. They also allow us to arrive at a simple implementable algorithm for computing localized homology. We implement this algorithm and show some experiments using simple cover constructions.

2 Background

In this section, we briefly discuss the algebraic tools required in our work. As it is infeasible to include a complete treatment, we sketch some of the basic ideas and include formal constructions only when necessary. For a more complete account, we refer to standard texts in the area [4, 5] and cite papers when needed. We organize this section as a continuity of ideas on capturing the topology of the spaces.

2.1 Topological Space

The fundamental object in topology is a *topological space*, an abstraction of a metric space. Rather than using a metric to define *open sets*, a topological space X is equipped with a set of open sets that define its connectivity. A subset $X_0 \subseteq X$ that is a topological space is a *subspace* and (X, X_0) is called a *pair*. A family $\mathcal{U} = \{X^i\}_i$ of subspaces $X^i \subseteq X$ is a *cover (covering)* of X if $X \subseteq \cup_i X^i$. We say that $\mathcal{U}$ covers X.

Suppose we have topological spaces X and Y and continuous maps $f\colon X \to Y$ and $g\colon Y \to X$ between them. If gf and fg are equal to the identity maps on the respective spaces, the spaces are *homeomorphic* and have the same *topological type*: $X \approx Y$. This is the most restrictive notion of equivalence in topology. We get a relaxation through the notion of *homotopy*. Given two maps $f_0, f_1\colon X \to Y$, if there is a continuous map $h\colon X \times [0,1] \to Y$ such that $h(x,0) = f_0(x)$ and $h(x,1) = f_1(x)$, then f_0 and f_1 are *homotopic* via *homotopy* h; $f_0 \simeq f_1$. Now, for our maps f and g above, if gf and fg are merely homotopic to the respective identities, then X and Y are *homotopy equivalent*: $X \simeq Y$.

For computation, we need a combinatorial structure for representing a topological space. Let $[n] = \{0, 1, \ldots, n\}$ be the first $n + 1$ natural numbers. An *n-simplex* σ is the convex hull of $n + 1$ affinely independent *vertices* $S = \{v^i\}_{i \in [n]}$ in $\mathbb{R}^d, d \geq n$. A simplex τ defined by $T \subseteq S$ is a *face* of σ. A *simplicial complex* K is a finite set of simplices that meet along faces, all of which are in K. A *subcomplex* of K is a subset $L \subseteq K$ that is also a simplicial complex. The *underlying space* $|K|$ of a simplicial complex K is $|K| = \cup_{\sigma \in K} \sigma$. A *triangulation* of a topological space X is a simplicial complex K such that $|K| \approx X$.

There is a standard realization for an n-simplex as follows. Let e^0 be the origin in $\mathbb{R}^n$ and $e^i = (0, \ldots, 1, \ldots, 0)$, $1 \leq i \leq n$, be the ith *standard basis vector* for $\mathbb{R}^n$ with a 1 in the ith position and 0's elsewhere. The *standard n-simplex* Δ^n is the convex hull of $\{e^i\}_{i \in [n]}$. The shaded triangle in Figure 7 is the standard 2-simplex. For any *indexing set* $J \subseteq [n]$, Δ^J is the face of Δ^n spanned by $\{e^j\}_{j \in J}$ with $\dim(\Delta^J) = \operatorname{card} J - 1$. Note that $\Delta^{[n]} = \Delta^n$. The standard simplex may be subdivided using the barycenters of its faces to produce the simplicial complex K^n with $|K^n| = \Delta^n$. Each non-empty face Δ^J of Δ^n has an associated vertex v^J in K^n. Δ^J is triangulated by subcomplex $K^J \subseteq K^n$ with

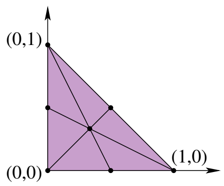

Figure 7. The shaded standard 2-simplex Δ^2 is the convex hull of the labeled vertices. Subdividing Δ^2 using the barycenters of the edges and the triangle gives the simplicial complex K^2 with underlying space Δ^2.

$|K^J| = \Delta^J$. Figure 7 displays this subdivision on the standard 2-simplex.

2.2 Homology

For a topological space X, the *homology groups* $H_n(X)$ are a family of Abelian groups for integers $n \geq 0$ with the following properties:

Functoriality: Each H_n is a *functor*, that is, for any continuous map $f\colon X \to Y$, there is an induced homomorphism $H_n(f)\colon H_n(X) \to H_n(Y)$, such that $H_n(fg) = H_n(f)H_n(g)$ and $H_n(i_X) = i_{H_n(X)}$, where i is the identity.

Homotopy Invariance: If $f, g\colon X \to Y$ are homotopic, then $H_n(f) = H_n(g)$. If f is a homotopy equivalence, then $H_n(f)$ is an isomorphism.

For any field F, there is a version of homology with coefficients in F that takes values in F-vector spaces. Throughout this paper, we always compute over a field. The rank of the nth vector space is the *nth Betti number* $\beta_n(X)$ of the space. In $\mathbb{R}^3$, the Betti numbers have intuitive meanings. β_0 measures the number of components of the complex. β_1 is the rank of a basis for the *tunnels*: loops that cannot be deformed to a point. β_2 counts the number of surfaces that enclose empty spaces.

Since we are interested in localizing homology, we need to understand the relationship between local and global homology of a space. The algebraic gadget that elucidates this relationship is the *Mayer-Vietoris sequence*. The Mayer-Vietoris sequence is does not give an algorithm, however. Fortunately, there is a geometric counter-part, the *Mayer-Vietoris blowup* which is defined as a subspace of a product space. For topological spaces X and Y, the *product space* $X \times Y$ is also a topological space. The connectivity of a product space is clearly related to the connectivity of its factors. We describe this relationship formally in the Appendix. For now, we note that we can obtain the homology of a product space from the homology of its factors.

2.3 Persistent Homology

As sketched in Section 1.3, we put together the blowup complex incrementally to see the local and global topologies at different times. For a topological space X, this construction gives a *filtration* $\{X^n\}_{n\geq 0}$, a nested sequence of subspaces: $\emptyset = X^0 \subseteq X^1 \cdots \subseteq X$. We call X a *filtered space*. We may similarly filter a simplicial complex to obtain a *filtered complex*.

Over fields, each space X^j has a kth homology group $H_k(X^j)$, a vector space whose rank $\beta_k(X^j)$ counts the number of topological attributes in dimension k. Viewing a filtration as a growing space, we see that topological attributes appear and cease. If we could track an attribute through the filtration, we could talk about its *lifetime* within it. The theory of *persistent homology* validates this intuition [9]. A filtration yields a *directed space*

$$\emptyset = X^0 \overset{i}{\hookrightarrow} \cdots \overset{i}{\hookrightarrow} X^j \cdots \overset{i}{\hookrightarrow} X,$$

where the maps i are the respective inclusions. Applying the kth dimensional homology functor H_k from Section 2.2 to both the spaces and the maps, we get another directed space

$$\emptyset = H_k\big(X^0\big) \xrightarrow{H_k(i)} \cdots \xrightarrow{H_k(i)} H_k\big(X^j\big) \xrightarrow{H_k(i)} \cdots \xrightarrow{H_k(i)} H_k(X)$$

where $H_k(i)$ are the respective induced maps. Persistent homology states that any directed vector space has a simple description. For an interval $[a, b]$, $a \in \mathbb{Z}^+, b \in \mathbb{Z}^+ \cup \{\infty\}$, let $F[a, b] = \{F_i\}_{n\geq 0}$ be a directed vector space over field F that is equivalent to F within the interval and empty elsewhere, and with identity transformations $F_i \to F_j$ within the interval. Under suitable finiteness hypotheses that are satisfied for all our spaces, the homology directed space may be written as a direct sum

$$\bigoplus_{i=0}^{s} F[a_i, b_i],$$

where the description is unique up to reordering of summands. In layman's terms, persistent homology states we may indeed track topological attributes and measure their lifetimes as intervals $[a_i, b_i]$. The *persistence barcode* is the finite multiset of lifetime intervals [2]. We can compute barcodes for arbitrary dimensional simplicial spaces over arbitrary fields using the persistence algorithm [9].

For this paper, we have extended the persistence algorithm to non-simplicial complexes, and we have enabled it to compute descriptions for non-boundary cycles. Sadly, the page limit prohibits further elaboration.

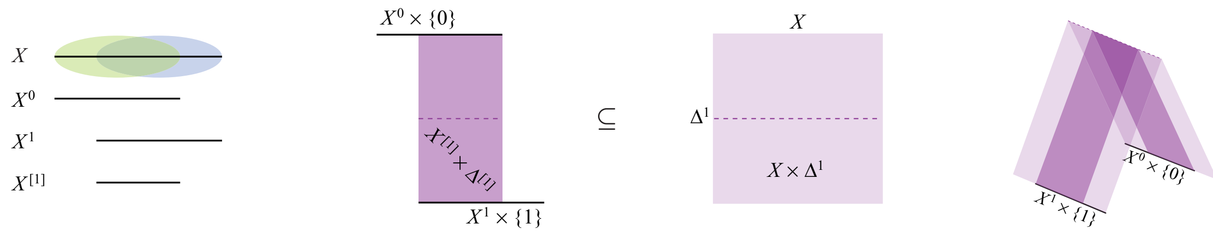

Figure 8. Blowup complex. Left: The cover $\mathcal{U} = \{X^0, X^1\}$ for space X also defines intersection $X^{\{0,1\}} = X^{[1]}$. Center: The blowup $X^{\mathcal{U}} \subseteq X \times \Delta^1$ is the union of three pieces shown. Here, Δ^1 is visualized as interval $[0, 1]$. Right: The function f on $X^{\mathcal{U}} \subseteq X \times \Delta^1$.

3 Localized Homology

In this section, we formalize the approach outlined in Section 1.3. We begin with a preliminary definition.

Definition 1 (localized homology) *Given a topological space X and a cover $\mathcal{U} = \{X^i\}_{i \in [n-1]}$, let $i\colon X^i \hookrightarrow X$ be inclusion, inducing $\iota_*\colon H_*\big(\dot{\bigcup}_{i\in[n-1]} X^i\big) \to H_*(X)$, where $\dot{\cup}$ is disjoint union. The* localized homology of X with respect to $\mathcal{U}$ *is the image of ι_*.*

In the rest of this section, we make this definition computational by deriving a filtered complex whose persistent homology is the localized homology. We choose a rich complex that carries more local information that the definition requires. We hope to examine the additional information in the near future. Our complex is the Mayer-Vietoris blowup complex $X^{\mathcal{U}}$, which we define and filter in Section 3.1. This complex enables us to refine Definition 1 using persistent homology.

Singular homology is not computational, however, as it deals with infinite-dimensional spaces. So, we adapt our definitions to simplicial spaces in Section 3.2. Given a simplicial complex X and a cover $\mathcal{U}$ of subcomplexes, we construct a simplicial blowup $X^{\mathcal{U}}$ that gives the same barcodes as the singular definition. While computable, the simplicial blowup is still a large complex. In Section 3.3, we describe a method for avoiding this construction by directly computing a small chain complex that yields equivalent barcodes. We then specify a natural basis and the boundary operator for this chain complex in Section 3.4. This specification allows us to construct a filtration directly from our space and cover, giving us a simple algorithm. Finally, in Section 3.5, we show localization results using an implementation of our algorithm.

3.1 Singular

Given an arbitrary topological space covered equipped with a cover, we blow up the space to incorporate the information contained in the cover: Each piece of the space expands according to the number of cover sets it falls within.

Definition 2 (Mayer-Vietoris blowup complex) *Given a topological space X with a cover $\mathcal{U} = \{X^i\}_{i\in[n-1]}$ of $n = \operatorname{card}\mathcal{U}$ sets, let $X^J = \cap_{j\in J} X^j$ for $J \subseteq [n-1]$. The* blowup complex $X^{\mathcal{U}} \subseteq X \times \Delta^{n-1}$ *of X and $\mathcal{U}$ is*

$$X^{\mathcal{U}} = \bigcup_{\emptyset \neq J \subseteq [n-1]} X^J \times \Delta^J. \tag{1}$$

$X^{\mathcal{U}}$ is equipped with two natural projection maps $\pi_X\colon X^{\mathcal{U}} \to X$ and $\pi_\Delta\colon X^{\mathcal{U}} \to \Delta^{n-1}$ given by the inclusion $X^{\mathcal{U}} \hookrightarrow X \times \Delta^{n-1}$ followed by projection onto the respective factors.

Example 1 (cover of two sets) Suppose X comes with cover $\mathcal{U} = \{X^0, X^1\}$ as shown on the left of Figure 8, where we represent X as an interval and draw ellipses to indicate the extent of the cover sets. The cover defines the intersection piece $X^{\{0,1\}} = X^{[1]}$. The blowup $X^{\mathcal{U}}$ is a subset of $X \times \Delta^1$ as shown the center of the figure, where we draw Δ^1 as the interval $[0, 1]$. Following Equation (1), $X^{\mathcal{U}}$ is the union of three pieces, corresponding to the three local regions the cover defines:

$$X^0 \times \Delta^{\{0\}} = X^0 \times \{0\},$$
$$X^1 \times \Delta^{\{1\}} = X^1 \times \{1\},$$
$$X^{[1]} \times \Delta^{[1]} = X^{[1]} \times [0, 1].$$

In constructing the blowup complex, we simply stretch certain pieces. Clearly then, the blowup complex has the same topology as the original space.

Proposition 1 (global) *The projection $\pi_X\colon X^{\mathcal{U}} \to X$ is a homotopy equivalence in the following cases:*

- *$\mathcal{U}$ is an open covering of a normal space, e.g. any subspace of $\mathbb{R}^n$,*

- $\mathcal{U}$ *is a covering of simplicial complexes by subcomplexes.*

Therefore, π_X *induces an isomorphism at the homology level. That is,* $X^{\mathcal{U}} \simeq X$ *and* $H_*(X^{\mathcal{U}}) \cong H_*(X)$ *[6].*

We now define a function f on $X^{\mathcal{U}}$ that assembles the pieces such that the persistent homology of the resulting filtration contains the localization solution. We first define a function on Δ^{n-1} by utilizing its triangulation K^{n-1}.

Definition 3 (height functions f, g) *For the face* $\emptyset \neq \Delta^J$ *of* Δ^{n-1}*, let* v^J *be the associated vertex in* K^{n-1}*. Define* $g\colon \Delta^{n-1} \to \mathbb{R}$ *linearly on the complex with* $g(v^J) = \operatorname{card} J - 1$*, and on* Δ^{n-1} *by identification. Define* $f\colon X^{\mathcal{U}} \to \mathbb{R}$ *by the composition* $X^{\mathcal{U}} \xrightarrow{\pi_\Delta} \Delta^{n-1} \xrightarrow{g} \mathbb{R}$.

We see f on the blowup complex on the right side of Figure 8. We filter the blowup complex using f.

Definition 4 (filtered blowup) *Let* $X_t^{\mathcal{U}} = f^{-1}([0,t])$*. The* filtered blowup complex *is the family* $\{X_t^{\mathcal{U}}\}_{t\geq 0}$.

In other words, when we visualize f as a height function on $X^{\mathcal{U}}$ as in the figures, $X_t^{\mathcal{U}}$ is everything in $X^{\mathcal{U}}$ below height t. At time 0, the blow up complex contains the local pieces of X. For Example 1, $X_0^{\mathcal{U}}$ is the two segments shown at the bottom of the hat shape in Figure 8.

Proposition 2 (local) *The space* $X_0^{\mathcal{U}}$ *is the disjoint union of the local pieces of the space, i.e.* $X_0^{\mathcal{U}} \approx \dot\cup_{i\in[n-1]} X^i$*. Therefore,*

$$H_k(X_0^{\mathcal{U}}) \cong \bigoplus_{i\in[n-1]} H_k(X^i),$$

that is, we get the homology of the local pieces at time 0.

So, we capture the local homology at time 0. At time $n-1$, the incremental construction is complete and $X_{n-1}^{\mathcal{U}} = X^{\mathcal{U}}$. Therefore, Proposition 1 asserts that $X_{n-1}^{\mathcal{U}}$ has the global homology of X. Applying persistent homology to the filtration, we get barcodes that describe the relationship between the local and global homology of the space. We may now state revise our definition for localization.

Definition 5 (localized homology) *Given a topological space* X *and a cover* $\mathcal{U} = \{X^i\}_{i\in[n-1]}$*, let* $i\colon X_0^{\mathcal{U}} \hookrightarrow X_{n-1}^{\mathcal{U}}$ *be inclusion, inducing* $\iota_*\colon H_*(X_0^{\mathcal{U}}) \to H_*(X_{n-1}^{\mathcal{U}})$*. The* localized homology of X with respect to $\mathcal{U}$ *is the image of* ι_*.

Equivalently, localized homology consists of the homology classes that exist at time 0 and continue to exist till time $n-1$. These classes correspond to persistence barcode intervals that contain both 0 and $n-1$.

3.2 Simplicial

The definitions in the last section assumed that X was a topological space, so the homology groups were all *singular* homology groups rather than those attached to a simplicial complex. Singular homology examines arbitrary maps of the standard simplex. The space of maps is an infinite dimensional space that is not computable. In this section, we modify our definitions for simplicial spaces and use simplicial homology which is finitely-generated and easily computable. We assume that we are given a simplicial complex X that represents a space of interest. We also restrict the cover $\mathcal{U}$ to consist of subcomplexes. Our task is twofold: we need to triangulate the blowup complex $X^{\mathcal{U}}$ and show that the simplicial homology of the resulting complex gives the same result as the singular method.

We begin by triangulating $X^{\mathcal{U}}$. Equation (1) states that $X^{\mathcal{U}}$ is a union of pieces of form $X^J \times \Delta^J$. Both terms $X^J = \cap_{j\in J} X^j$ and Δ^J are simplicial, giving us a product of simplicial complexes. Given total orderings on the vertex sets, there is a canonical way to triangulate a product space that we omit due to lack of space. This triangulation gives us the simplicial blowup complex. We now define both the simplicial blowup and its filtration at once.

Definition 6 (filtered simplicial blowup) *Let* X *be a simplicial complex and* $\mathcal{U} = \{X^i\}_{i\in[n-1]}$ *be a cover of* n *subcomplexes. For* $J \subseteq [n-1]$*, let* $X^J = \cap_{j\in J} X^j$ *and*

$$X_t^{\mathcal{U}} = \bigcup_{\substack{J\subseteq[n-1] \\ 0<\operatorname{card} J \leq t+1}} X^J \times \Delta^J. \tag{2}$$

The blowup complex of X and $\mathcal{U}$ *is* $X^{\mathcal{U}} = X_{n-1}^{\mathcal{U}} \subseteq X \times \Delta^{n-1}$ *with projections* $\pi_X\colon X^{\mathcal{U}} \to X$ *and* $\pi_\Delta\colon X^{\mathcal{U}} \to \Delta^{n-1}$*. The* filtered blowup complex *is the family* $\{X_t^{\mathcal{U}}\}_{t\geq 0}$.

Note that $X_t^{\mathcal{U}}$ is defined for all $t \in \mathbb{R}, t \geq 0$, but the complex changes only at integer values. Figure 9 constructs the blowup complex for Example 1 in simplicial form. Definition 6 mimics the singular definitions 3 and 4. For example, the piece $X^{[1]} \times [0,1]$ is completed at time 1 in the hat shape in Figure 8, and the triangulation of the corresponding piece $X^{[1]} \times \Delta^{[1]}$ in Figure 9 also arrives at time 1.

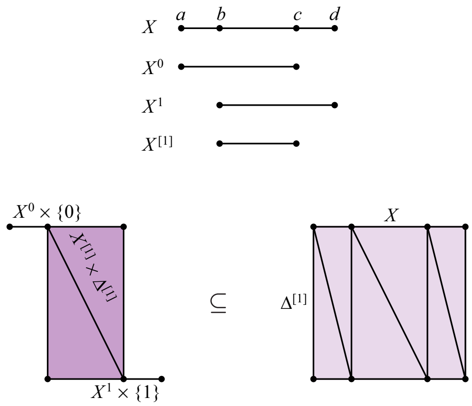

Figure 9. Top: The simplicial cover $\mathcal{U} = \{X^0, X^1\} = \{ac, bd\}$ for simplicial complex X defines $X^{\{0,1\}} = X^{[1]} = bc$. Bottom: The simplicial blowup complex $X^{\mathcal{U}} \subseteq X \times \Delta^1$.

To complete our task, we need to show that the new simplicial definition has the same structure as the singular one from the last section. The underlying space $|X|$ of X is a topological space with the cover $|\mathcal{U}| = \{|X^i|\}_{i\in[n-1]}$. Applying Definition 2, we get the singular blowup complex $|X|^{|\mathcal{U}|}$ that looks like the blowup in Figure 8. Clearly, the blowups are identical at integer values for t. For non-integer t, the two are homotopic since topology only changes at integer values.

Proposition 3 *Given a simplicial complex X and a cover $\mathcal{U} = \{X^i\}_{i\in[n-1]}$ of subcomplexes, let $|\mathcal{U}| = \{|X^i|\}_{i\in[n-1]}$. There exists a canonical homeomorphism $\varphi\colon |X^{\mathcal{U}}| \to |X|^{|\mathcal{U}|}$ with restriction $\varphi_t\colon |X_t^{\mathcal{U}}| \to |X|^{|\mathcal{U}|}$ such that:*

1. *For $t \geq 0$, $\varphi_t\left(|X_t^{\mathcal{U}}|\right) \subseteq |X|_t^{|\mathcal{U}|}$.*
2. *For $t \in [n-1]$, φ_t is a homeomorphism onto its image.*
3. *For $t \geq 0$, φ_t is a homotopy equivalence.*

Corollary 1 *The map φ of filtered spaces in Proposition 3 induces an isomorphism of directed Abelian groups from $H_k(|X_t^{\mathcal{U}}|)$ to $H_k(|X|_t^{|\mathcal{U}|})$ for non-negative integers k and $t \geq 0$. So, their barcodes are equivalent.*

Instead of using the singular definition in the last section, we use the simplicial definition to compute the barcodes of the blowup complex. The former definition was not computable; the latter is.

3.3 Chain Complex

While the simplicial definition of the blowup complex is computable, it is not efficient as we have to build the blowup complex and triangulate all its pieces. In this section, we eliminate the triangulation step and work directly at the chain level. Rather than computing homology using the chain complex attached to the simplicial blowup complex, we utilize a smaller chain complex that gives equivalent barcodes and is computed directly from the complex and the cover. This is essentially a version of the cellular complex.

We begin by examining the chain complex attached to the simplicial blowup complex. It follows directly from Equation (2) that the filtered chain complex for the blowup complex is the family $\{C_*(X^{\mathcal{U}})_t\}_{t\geq 0}$, where $C_*(X^{\mathcal{U}})_t \subseteq C_*(X \times \Delta^{n-1})$ is

$$C_*\left(X^{\mathcal{U}}\right)_t = \sum_{\substack{J\subseteq[n-1] \\ 0\,<\,\mathrm{card}\,J\,\leq\,t+1}} C_*\left(X^J \times \Delta^J\right), \qquad (3)$$

and $C_*(X^{\mathcal{U}}) = C_*(X^{\mathcal{U}})_{n-1}$. For each piece $C_*(X^J \times \Delta^J)$, we triangulated $X^J \times \Delta^J$ in the previous section. To avoid triangulating the product, we define a smaller chain complex.

Definition 7 (filtered blowup chain complex) *Let X be a simplicial complex and $\mathcal{U} = \{X^i\}_{i\in[n-1]}$ be a cover of n subcomplexes. For $J \subseteq [n-1]$, let $X^J = \cap_{j\in J} X^j$ and $C_*^{\mathcal{U}}(X)_t \subseteq C_*(X) \otimes C_*(\Delta^{n-1})$ be*

$$C_*^{\mathcal{U}}(X)_t = \sum_{\substack{J\subseteq[n-1] \\ 0\,<\,\mathrm{card}\,J\,\leq\,t+1}} C_*\left(X^J\right) \otimes C_*\left(\Delta^J\right). \qquad (4)$$

The blowup chain complex of X and $\mathcal{U}$ *is* $C_*^{\mathcal{U}}(X) = C_*^{\mathcal{U}}(X)_{n-1}$. *The* filtered blowup chain complex *is the family* $\{C_*^{\mathcal{U}}(X)_t\}_{t\geq 0}$.

The two filtered complexes defined by Equations (3) and (4) give the same localized homology.

Proposition 4 *Given simplicial space X and simplicial cover $\mathcal{U} = \{X^i\}_{i\in[n-1]}$, there is a chain map $\mathcal{A}\colon C_*(X^{\mathcal{U}}) \to C_*^{\mathcal{U}}(X)$ that induces an isomorphism of directed Abelian groups from $H_k(C_*(X^{\mathcal{U}})_t)$ to $H_k(C_*^{\mathcal{U}}(X)_t)$ for non-negative integers k and $t \geq 0$. Consequently, the barcodes of the two chain complexes are equivalent.*

The dearth of space relegates the definition of $\mathcal{A}$ and the proof of the proposition to the Appendix.

3.4 Algorithm

In the previous section, we showed that the chain complex $C_*^{\mathcal{U}}(X)$ has the same localized homology as the simplicial blowup complex. To compute the homology, we need a basis for $C^{\mathfrak{U}}_*(X)$ and the boundary homomorphism. By Equation (4), $C^{\mathfrak{U}}_*(X)$ is a sum of terms of the form $C_*(X^J) \otimes C_*(\Delta^J)$ over nonempty sets $J \subseteq [n-1]$. As both terms of the product are simplicial, they give rise to a canonical basis.

Proposition 5 (basis) *A basis for $C^{\mathfrak{U}}_k(X)$ is the set composed of elements $\sigma \otimes \Delta^J$ for all $\emptyset \neq J \subseteq [n-1]$ and simplices $\sigma \in X^J$ where $\dim \sigma + \dim \Delta^J = k$.*

To define the boundary, we impose total orderings on the vertices of X and Δ^{n-1}.

Proposition 6 (boundary homomorphism) *Let $\sigma \otimes \Delta^J$ be a basis element for $C^{\mathfrak{U}}_k(X)$. Then,*

$$\partial\big(\sigma \otimes \Delta^J\big) = \partial\sigma \otimes \Delta^J + (-1)^{\dim \sigma} \sigma \otimes \partial\Delta^J. \quad (5)$$

Example 2 Consider our simplicial example in Figure 9. We no longer triangulate the blowup complex. Its single 2-cell is represented by the basis element $bc \otimes \Delta^{[1]} = bc \otimes [0,1]$ for $C^{\mathfrak{U}}_2(X)$. The boundary of this cell is composed of the four 1-cells, represented as tensor products:

$$\partial\,(bc \otimes [0,1]) = c \otimes [0,1] - b \otimes [0,1] - bc \otimes \{1\} + bc \otimes \{0\}.$$

The first two terms originate from the first sum in Equation (5) and correspond to the two sides. The second two originate from the second sum and represent the two copies of bc on the top and the bottom of the cell, respectively, oriented oppositely.

In practice, we represent basis element $\sigma \otimes \Delta^J$ as the pair (σ, J) and use Equation (5) as the boundary operator. This representation feeds directly into the persistence algorithm, giving us the barcode, localization through Definition 5, and cycle descriptions. We now have a simple algorithm for computing localized homology.

In general, we also do not need to compute the entire blowup complex. Computing the kth homology group requires the $(k+1)$-*skeleton*, cells with dimension less than or equal to $k+1$. Since only the first d homology groups of a d-dimensional space may be nontrivial, the largest complex we build is the $(d+1)$-skeleton, which is generally much smaller than the full blowup.

3.5 Experiments

We have implemented our algorithm as part of a library of programs for algebraic topology. We use the same code for computing homology and persistent homology of a simplicial complex and that of its blowup complex, represented as a abstract complex of simplex products. Since our focus here is on demonstrating our method, we only consider two naive methods for cover generation based on *random ϵ-balls* and *tilings*.

The space in Figure 10(a) is a 2-dimensional punctured sheet embedded in $\mathbb{R}^2$. The complex has 3,360 simplices. Computing homology, we get the basis 1-cycles shown in Figure 10(b), where one cycle goes around two holes. We use a disc of radius 10% the diameter of the space as our local element and take the closure of simplices that fall within a disc as a set in the cover. Figure 10(a) highlights the 54th set, the first set that contains the medium-sized cycle. Our method covers the complex with 119 sets, colored transparently in Figure 10(c) and the 2-skeleton of resulting the blowup complex has 101,402 cells. In Figure 10(d), we show the projection of the blowup 1-cycles to the base space: our random cover localizes the two smaller holes but the not the largest since it is not contained within a single set in the cover.

We next detect a two-dimensional cycle (void) carved out from the center of the cubical block shown in Figure 11. We cover the complex of 67,370 simplices using a systematic method based on tiling Euclidean spaces. The method generates covers of size at most $2^d + 1$ for complexes embedded in $\mathbb{R}^d$ (i.e. 9 in $\mathbb{R}^3$) and guarantees localization of cycles half the size of a tile. The 3-skeleton of the blowup complex has 474,416 cells and localizes the void.

4 Conclusion

In this paper, we introduce the theory of localized homology. The key contribution of our approach is combining the classic blowup construction with the recent technique of persistent homology. We ground our method in the general setting of arbitrary spaces and singular homology. As such, our theory does not consist of ad-hoc techniques that only work in particular dimensions and spaces, such as finding tunnels on surfaces, but in arbitrary dimensions and spaces, such as the void in Figure 11. We know of no other comparable work. We provide equivalent definitions for simplicial spaces to arrive at an efficient method that utilizes a smaller chain complex. Finally, we implement our method and give results.

The major issue we do not address is cover construction. Our method is very flexible as it does not place any restrictions on the geometry or topology of the cover sets, as seen in the void example. We can tailor the cover construction to the localization requirements in different settings and even utilize multiple schemes in tandem. Having constructed a robust localization engine, we plan to examine cover construction next.

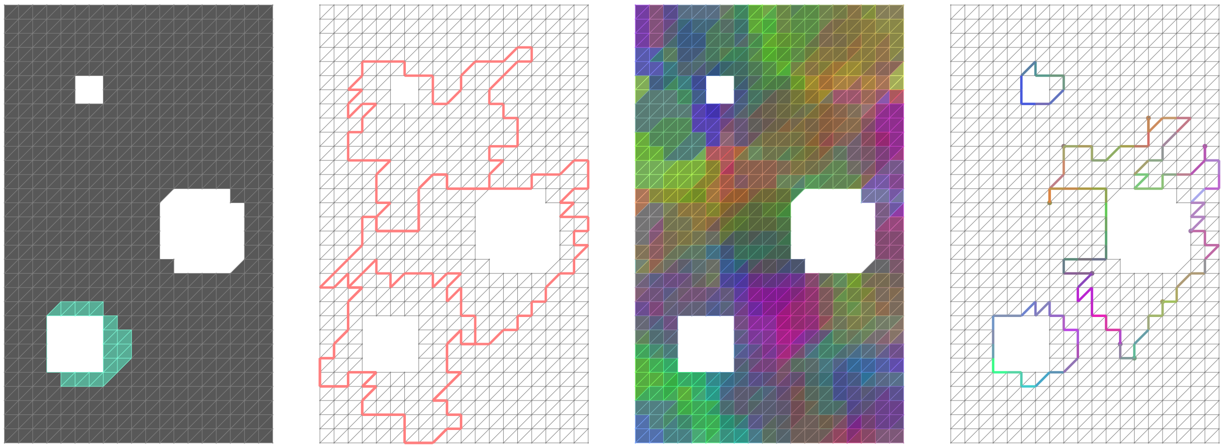

(a) Complex and cover set (b) Homology cycles (c) ϵ-disc cover (d) Blowup cycles

Figure 10. Punctured sheet. The complex (a) highlights the first set in the cover that contains the medium hole. The 1-cycles computed with homology (b) are nonlocal and one goes around multiple holes. The colored transparent sets (c), based on ϵ-discs, cover the complex. We project the 1-cycles of the blowup complex (d) to get a localized description for the two smaller holes.

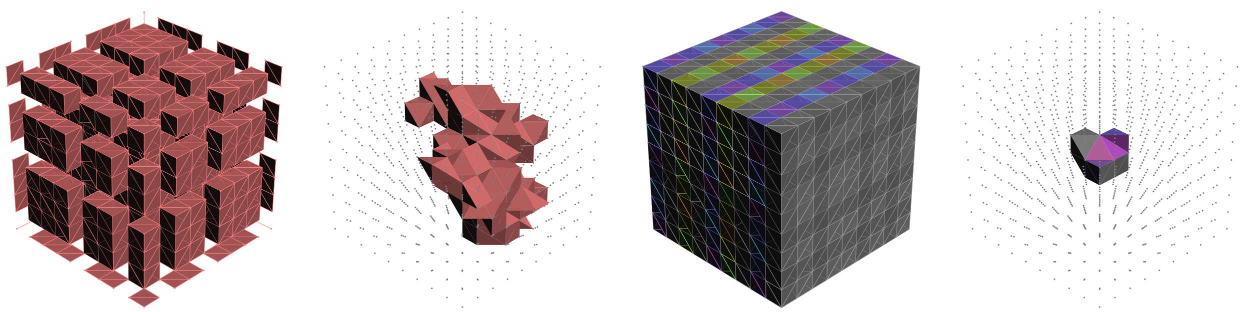

Figure 11. A void. We carve out the center of a cubical block. From left: the first set of a cover, the large 2-cycle computed with homology, the complex covered by 8 sets, and the much smaller 2-cycle computed with localized homology.

References

[1] APAYDIN, M. S., BRUTLAG, D. L., GUESTRIN, C., HSU, D., LATOMBE, J.-C., AND VARMA, C. Stochastic roadmap simulation: An efficient representation and algorithm for analyzing molecular motion. *Journal of Computational Biology 10* (2003), 257–281.

[2] CARLSSON, G., ZOMORODIAN, A., COLLINS, A., AND GUIBAS, L. J. Persistence barcodes for shapes. In *Proceedings of the Symposium on Geometry Processing* (2004), pp. 127–138.

[3] FANG, Q., GAO, J., AND GUIBAS, L. J. Locating and bypassing routing holes in sensor networks. In *IEEE INFOCOM* (2004), pp. 2458–2468.

[4] GREENBERG, M. J., AND HARPER, J. R. *Algebraic Topology: A First Course*, vol. 58 of *Mathematics Lecture Note Series*. Benjamin/Cummings Publishing Co., Reading, MA, 1981.

[5] HATCHER, A. *Algebraic Topology*. Cambridge University Press, Cambridge, UK, 2001.

[6] SEGAL, G. B. Equivariant K-theory. *Inst. Hautes Etudes Sci. Publ. Math. 34* (1968), 129–151.

[7] WOOD, Z., HOPPE, H., DESBRUN, M., AND SCHRÖDER, P. Removing excess topology from isosurfaces. *ACM Transactions on Graphics 23*, 2 (Apr. 2004), 190–208.

[8] ZOMORODIAN, A. *Topology for Computing*. Cambridge University Press, New York, NY, 2005.

[9] ZOMORODIAN, A., AND CARLSSON, G. Computing persistent homology. *Discrete and Computational Geometry 33*, 2 (2005), 249–274.

Appendix

The Alexander-Whitney map describes the relationship of a product space to its factors.

Theorem 1 (Alexander-Whitney) *Let X and Y be simplicial complexes with vertex orderings $<_X$ and $<_Y$ defining the product simplicial complex $X \times Y$. For maps of simplicial complexes that preserve the vertex orderings, there are natural chain maps*

$$C_*(X \times Y) \underset{\mathcal{S}}{\overset{\mathcal{A}}{\rightleftarrows}} C_*(X) \otimes C_*(Y)\,, \tag{6}$$

that induce isomorphisms on homology groups. Here, $\mathcal{A}$ is the Alexander-Whitney *map and $\mathcal{S}$ is the* shuffle *homomorphism. For pairs (X, X_0) and (Y, Y_0), we get a relative version*

$$C_*(X \times Y,\ (X \times Y_0) \cup (X_0 \times Y)) \underset{\mathcal{S}}{\overset{\mathcal{A}}{\rightleftarrows}} C_*(X, X_0) \otimes C_*(Y, Y_0)\,. \tag{7}$$

We now relate the two filtered complexes defined by Equations (3) and (4) in Section 3.3. Observe that both complexes sum over the same variable J, and the summand of the first complex $C_*(X^J \times \Delta^J)$ maps naturally via the Alexander-Whitney to the summand of the second complex $C_*(X^J) \otimes C_*(\Delta^J)$, according to Equation (6).

Definition 8 *Let $\mathcal{A}\colon C_*(X^{\mathcal{U}}) \to C_*^{\mathcal{U}}(X)$ be the map whose the restriction to any summand $C_*(X^J \times \Delta^J)$ in Equation* (3) *is the Alexander-Whitney map with values in $C_*(X^J) \otimes C_*(\Delta^J)$.*

There is at most one map $\mathcal{A}$ that satisfies the requirements as we specify the map on a generating family of subcomplexes. The existence $\mathcal{A}$ is guaranteed by two facts: the naturality of the Alexander-Whitney map and $C_*(X_0 \cap X_1) = C_*(X_0) \cap C_*(X_1)$ for subcomplexes $X_0, X_1 \subseteq X$. Clearly, $\mathcal{A}(C_*(X^{\mathcal{U}})_t) \subseteq C_*^{\mathcal{U}}(X)_t$, so $\mathcal{A}$ is a chain map. We use this chain map to restate and prove Proposition 4.

Proposition 4 *Given simplicial space X and simplicial cover $\mathcal{U} = \{X^i\}_{i \in [n-1]}$, the chain map $\mathcal{A}\colon C_*(X^{\mathcal{U}}) \to C_*^{\mathcal{U}}(X)$ induces an isomorphism of directed Abelian groups from $H_k(C_*(X^{\mathcal{U}})_t)$ to $H_k(C_*^{\mathcal{U}}(X)_t)$ for non-negative integers k and $t \geq 0$. Consequently, the barcodes of the two chain complexes are equivalent.*

Proof: It suffices to show that the restriction $\mathcal{A}_t\colon C_*(X^{\mathcal{U}})_t \to C_*^{\mathcal{U}}(X)_t$ induces an isomorphism on homology groups on the integer values of t where topological changes occur. By inductive use of the long exact homology sequence associated with the short exact sequence of chain complexes, it suffices to prove that the following induced map on subquotients induces an isomorphism on homology groups:

$$\hat{\mathcal{A}}_t\colon C_*\big(X^{\mathcal{U}}\big)_t / C_*\big(X^{\mathcal{U}}\big)_{t-1} \to C_*^{\mathcal{U}}(X)_t / C_*^{\mathcal{U}}(X)_{t-1}\,.$$

From Equation (3) we have a direct sum decomposition for the domain of $\hat{\mathcal{A}}_t$:

$$C_*\big(X^{\mathcal{U}}\big)_t / C_*\big(X^{\mathcal{U}}\big)_{t-1} \cong \bigoplus_{\substack{J \subseteq [n-1] \\ \operatorname{card} J = t+1}} C_*\big(X^J \times \Delta^J,\ X^J \times \partial\Delta^J\big)\,, \tag{8}$$

where ∂ is the boundary operator. Similarly, from Equation (4) we have a direct sum decomposition for the codomain of $\hat{\mathcal{A}}_t$:

$$C_*^{\mathcal{U}}(X)_t / C_*^{\mathcal{U}}(X)_{t-1} \cong \bigoplus_{\substack{J \subseteq [n-1] \\ \operatorname{card} J = t+1}} C_*\big(X^J\big) \otimes C_*\big(\Delta^J,\ \partial\Delta^J\big)\,. \tag{9}$$

The restriction of $\hat{\mathcal{A}}_t$ to each summand in Equation (8) maps the summand to the corresponding summand in Equation (9),

$$C_*\big(X^J \times \Delta^J, X^J \times \partial\Delta^J\big) \xrightarrow{\hat{\mathcal{A}}_t} C_*\big(X^J\big) \otimes C_*\big(\Delta^J, \partial\Delta^J\big)\,,$$

as the restriction is the Alexander-Whitney map obtained by setting $X = X^J$, $Y = \Delta^J$, $X_0 = \emptyset$, and $Y_0 = \partial\Delta^J$ in Equation (7). Therefore, the restriction induces an isomorphism on homology groups and the theorem follows. $\square$