

Computing Minimal Persistent Cycles: Polynomial and Hard Cases*

Tamal K. Dey[†]

Tao Hou[‡]

Sayan Mandal[§]

Abstract

Persistent cycles, especially the minimal ones, are useful geometric features functioning as augmentations for the intervals in a purely topological persistence diagram (also termed as barcode). In our earlier work, we showed that computing minimal 1-dimensional persistent cycles (persistent 1-cycles) for finite intervals is NP-hard while the same for infinite intervals is polynomially tractable. In this paper, we address this problem for general dimensions with $\mathbb{Z}_2$ coefficients. In addition to proving that it is NP-hard to compute minimal persistent d -cycles ($d > 1$) for both types of intervals given arbitrary simplicial complexes, we identify two interesting cases which are polynomially tractable. These two cases assume the complex to be a certain generalization of manifolds which we term as *weak pseudomanifolds*. For finite intervals from the d_{th} persistence diagram of a weak $(d + 1)$ -pseudomanifold, we utilize the fact that persistent cycles of such intervals are null-homologous and reduce the problem to a minimal cut problem. Since the same problem for infinite intervals is NP-hard, we further assume the weak $(d + 1)$ -pseudomanifold to be embedded in $\mathbb{R}^{d+1}$ so that the complex has a natural dual graph structure and the problem reduces to a minimal cut problem. Experiments with both algorithms on scientific data indicate that the minimal persistent cycles capture various significant features of the data.

*Supported by NSF grants CCF-1740761 and CCF-1839252.

[†]Department of Computer Science and Engineering, The Ohio State University. dey.8@osu.edu

[‡]Department of Computer Science and Engineering, The Ohio State University. hou.332@osu.edu

[§]Department of Computer Science and Engineering, The Ohio State University. mandal.25@osu.edu

1 Introduction

Persistent homology [15], which captures essential topological features of data, has proven to be a useful stable descriptor since Edelsbrunner et al. [16] first proposed the algorithm for its computation. The understanding of topological persistence was later expanded by several works [5, 9, 11, 31] in terms of both theory and computation. To make use of persistent homology, one typically computes a *persistence diagram* (also called *barcode*) which is a set of intervals with birth and death points. Besides just utilizing the set of intervals, some applications [13, 30] need persistence diagrams augmented with representative cycles for the intervals for gaining more insight into the data. These representative cycles, termed as *persistent cycles* [13], have been studied by Wu et al. [30], Obayashi [24], and Dey et al. [13] recently from the view-point of optimality.

Although the original persistence algorithm of Edelsbrunner et al. [16] implicitly computes persistent cycles, it does not necessarily provide minimal ones. In an earlier work [13], we showed that it is NP-hard to compute minimal persistent 1-cycles (cycles for 1-dimensional homology groups) when the given interval is finite. Interestingly, the same for infinite intervals turned out to be computable in polynomial time [13]. This naturally leads to the following questions: Are there other interesting cases beyond 1-dimension for which minimal persistent cycles can be computed in polynomial time? Also, what are the cases that are NP-hard? In this paper, we settle the complexity question for computing minimal persistent cycles with $\mathbb{Z}_2$ coefficients in general dimensions. We first show that when $d \geq 2$, computing minimal persistent d -cycles for both finite and infinite intervals is NP-hard in general. We then identify a special but important class of simplicial complexes, which we term as *weak $(d+1)$ -pseudomanifolds*, whose minimal persistent d -cycles can be computed in polynomial time. A weak $(d+1)$ -pseudomanifold* is a generalization of a $(d+1)$ -manifold and is defined as follows:

Definition 1.1. A simplicial complex K is a *weak $(d+1)$ -pseudomanifold* if each d -simplex is a face of no more than two $(d+1)$ -simplices in K .

Specifically, we find that if the given complex is a weak $(d+1)$ -pseudomanifold, the problem of computing minimal persistent d -cycles for finite intervals can be cast into a minimal cut problem (see Section 3) due to the fact that persistent cycles of such kind are null-homologous in the complex. However, when $d \geq 2$ and intervals are infinite, the computation of the same becomes NP-hard (see Section 5). Nonetheless, for infinite intervals, if we assume that the weak $(d+1)$ -pseudomanifold is embedded in $\mathbb{R}^{d+1}$, the minimal persistent cycle problem reduces to a minimal cut problem (see Section 4) and hence belongs to P. Note that a simplicial complex embedded in $\mathbb{R}^{d+1}$ is automatically a weak $(d+1)$ -pseudomanifold. Also note that while there is an algorithm [8] in the non-persistence setting which computes minimal d -cycles by minimal cuts, the non-persistence algorithm assumes the $(d+1)$ -complex to be embedded in $\mathbb{R}^{d+1}$. Our algorithm for finite intervals, to the contrary, does not need the embedding assumption.

In order to make our statements about the hardness results precise, we let PCYC-FIN $_d$ denote the problem of computing minimal persistent d -cycles for finite intervals when the given simplicial complex is arbitrary, and let PCYC-INF $_d$ denote the same problem for infinite intervals (see definitions of Problem 2.1 and 2.2). We also let WPCYC-FIN $_d$ denote a subproblem[†] of PCYC-FIN $_d$ and let WPCYC-INF $_d$, WEPCYC-INF $_d$ denote two subproblems of PCYC-INF $_d$, with the subproblems requiring additional constraints on the given simplicial complex. Table 1 lists the

*The naming of weak pseudomanifold is adapted from the commonly accepted name *pseudomanifold* (see Definition A.1).

[†]For two problems P_1 and P_2 , P_2 is a *subproblem* of P_1 if any instance of P_2 is an instance of P_1 and P_2 asks for computing the same solutions as P_1 .

hardness results for all problems of interest, where the column “Restriction on K ” specifies the additional constraints subproblems require on the given simplicial complex K . Note that WPCYC-INF $_d$ being NP-hard trivially implies that PCYC-INF $_d$ is NP-hard.

Table 1: Hardness results for minimal persistent cycle problems with bold results denoting new findings.

<i>Problem</i>	<i>Restriction on K</i>	<i>d</i>	<i>Hardness</i>
PCYC-FIN $_d$	—	≥ 1	NP-hard
WPCYC-FIN $_d$	K a weak $(d + 1)$ -pseudomanifold	≥ 1	Polynomial
PCYC-INF $_d$	—	$= 1$	Polynomial
WPCYC-INF $_d$	K a weak $(d + 1)$ -pseudomanifold	≥ 2	NP-hard
WEPCYC-INF $_d$	K a weak $(d + 1)$ -pseudomanifold in $\mathbb{R}^{d+1}$	≥ 2	Polynomial

Main contributions. We summarize our contributions as follows:

- We prove the NP-hardness of PCYC-FIN $_d$ and WPCYC-INF $_d$ for all $d \geq 2$.
- We present two polynomial time algorithms for WPCYC-FIN $_d$ and WEPCYC-INF $_d$ when $d \geq 1$, based on the duality of minimal persistent cycles and minimal cuts. Other than the minimal cut computation, steps in both algorithms run in linear or almost linear time.

1.1 Related works

In the context of computing optimal cycles, most works have been done in the non-persistence setting. These works compute minimal cycles for homology groups of a given simplicial complex. Only very few works address the problem while taking into account the persistence. We review some of the relevant works below.

Minimal cycles for homology groups. In terms of computing minimal cycles for homology groups, two problems are of most interest: the localization problem and the minimal basis problem. The localization problem asks for computing a minimal cycle in a homology class and the minimal basis problem asks for computing a set of generating cycles for a homology group whose sum of weights is minimal. With $\mathbb{Z}_2$ coefficients, these two problems are in general hard. Specifically, Chambers et al. [4] proved that the localization problem over dimension one is NP-hard when the given simplicial complex is a 2-manifold. Chen and Freedman [8] proved that the localization problem is NP-hard to approximate with fixed ratio over arbitrary dimension. They also showed that the minimal basis problem is NP-hard to approximate with fixed ratio over dimension greater than one. For one-dimensional homology, Dey et al. [14] proposed a polynomial time algorithm for the minimal basis problem. Several other works [3, 7, 12, 18] address variants of the two problems while considering special input classes, alternative cycle measures, or coefficients for homology other than $\mathbb{Z}_2$.

In this work, we use graph cuts and their duality extensively. The duality of cuts on a planar graph and separating cycles on the dual graph has long been utilized to efficiently compute maximal flows and minimal cuts on planar graphs, a topic for which Chambers et al. [4] provide a comprehensive review. In their paper [4], Chambers et al. discover the duality between minimal cuts of a surface-embedded graph and minimal homologous cycles in a dual complex, and then devise $O(n \log n)$ algorithms for both problems assuming the genus of the surface to be fixed. Chen and Freedman [8] proposed an algorithm which computes a minimal non-bounding d -cycle given a $(d + 1)$ -complex

embedded in $\mathbb{R}^{d+1}$, utilizing a natural duality of d -cycles in the complex and cuts in the dual graph. The minimal non-bounding cycle algorithm can be further extended to solve the localization problem and the minimal basis problem over dimension d given a $(d + 1)$ -complex embedded in $\mathbb{R}^{d+1}$.

Persistent cycle. As pointed out earlier, our main focus is the optimality of representative cycles in the persistence framework. Some early works [17, 19] address the representative cycle problem for persistence by computing minimal cycles at the birth points of intervals without considering what actually die at the death points. Wu et al. [30] proposed an algorithm computing minimal persistent 1-cycles for finite intervals using an annotation technique and heuristic search. However, the time complexity of the algorithm is exponential in the worst-case. Obayashi [24] casts the minimal persistent cycle problem for finite intervals into an integer program, but the rounded result of the relaxed linear program is not guaranteed to be optimal. Dey et al. [13] formalizes the definition of persistent cycles for both finite and infinite intervals. They also proved the NP-hardness of computing minimal persistent 1-cycles for finite intervals and proposed a polynomial time algorithm for computing non-optimal ones which are still good in practice.

2 Preliminaries

In this section we present some concepts necessary for presenting the results in this paper.

Simplicial complex. A *simplicial complex* K is a collection of *simplices* which are abstractly defined as subsets of a ground set called the *vertex set* of K . If a simplex σ is in K , then all its subsets called its *faces* are also in K . The simplex σ is also referred to as a q -simplex if the cardinality of the vertex set of σ is $q + 1$. A q -*face* of σ is a q -simplex being a face of σ and a q -*coface* of σ is a q -simplex having σ as a face. We call a q -simplex of K a *boundary q -simplex* if it has less than two $(q + 1)$ -cofaces in K . A *simplicial set* is a set of simplices and the *closure* of a simplicial set Σ is the simplicial complex consisting of all the faces of the simplices in Σ . A simplicial complex is *finite* if it contains finitely many simplices. In this paper, we only consider finite simplicial complexes.

If each vertex of a simplicial complex K is a point in a Euclidean space, then each simplex of K can be interpreted as the convex hull of its vertices. The simplicial complex K is said to be *embedded in* the Euclidean space if the interiors of all its simplices are disjoint. The *underlying space* of K , denoted by $|K|$, is the point-wise union of all the simplices of K .

Definition 2.1 (Oriented simplex [23]). A q -simplex with an ordering of its vertices is an *oriented q -simplex*. For each q -simplex σ ($q > 0$), there are exactly two equivalent classes of vertex orderings, resulting in two oriented q -simplices of σ . We refer to them as the *oppositely oriented q -simplices*.

Remark 2.1. Any simplex by default is unoriented. We denote an unoriented q -simplex σ spanned by vertices $v_0, \dots, v_q$ as $\sigma = \{v_0, \dots, v_q\}$ and an oriented q -simplex $\vec{\sigma}$ as $\vec{\sigma} = [v_0, \dots, v_q]$, where $v_0, \dots, v_q$ specify the ordering of the spanning vertices.

Filtration. A *filtration* $\mathcal{F}$ of a simplicial complex K is a filtered sequence of subcomplexes of K , $\mathcal{F} : \emptyset = K_0 \subseteq K_1 \subseteq \dots \subseteq K_n = K$, such that K_i and K_{i-1} differ by one simplex denoted by $\sigma_i^{\mathcal{F}}$. We let i be the *index* of $\sigma_i^{\mathcal{F}}$ in $\mathcal{F}$ and denote it as $\text{ind}(\sigma_i^{\mathcal{F}}) = i$. A subcomplex K_i in the filtered sequence of $\mathcal{F}$ is also referred to as a *partial complex*.

Simplicial homology. We provide a brief overview of simplicial homology used in this paper. See any standard book on the topic, e.g. [23]. Let $q \geq 0$, K be a simplicial complex, and $\mathbb{G}$ be an abelian group. The q th chain group $C_q(K; \mathbb{G})$ is defined to be the abelian group containing all finite sums of the form $\sum_i n_i \vec{\sigma}_i$, where $n_i \in \mathbb{G}$ and $\vec{\sigma}_i$ is an oriented q -simplex of K . Each element in $C_q(K; \mathbb{G})$ is called a q -chain of K . Note that for two oppositely oriented q -simplices $\vec{\sigma}$ and $\vec{\sigma}'$, we have that $n\vec{\sigma} = (-n)\vec{\sigma}'$ for any $n \in \mathbb{G}$. Therefore, $C_q(K; \mathbb{G})$ can be interpreted as a direct sum of N_q copies of $\mathbb{G}$ where N_q is the number of q -simplices of K and each copy of $\mathbb{G}$ corresponds to a q -simplex of K . The q th boundary operator $\partial_q : C_q(K; \mathbb{G}) \rightarrow C_{q-1}(K; \mathbb{G})$ is a group homomorphism such that for any oriented q -simplex $[v_0, \dots, v_q]$

$$\partial_q([v_0, \dots, v_q]) = \sum_{i=0}^q (-1)^i [v_0, \dots, \widehat{v}_i, \dots, v_q]$$

where the notation $[v_0, \dots, \widehat{v}_i, \dots, v_q]$ means that $\widehat{v}_i$ is deleted from the simplex. For brevity, we often omit the subscript of the boundary operator ∂_q and denote it as ∂ when this does not cause any confusion. The kernel of ∂_q is called the q th cycle group of K and is denoted as $Z_q(K; \mathbb{G})$. The image of ∂_{q+1} is called the q th boundary group of K and is denoted as $B_q(K; \mathbb{G})$. A q -chain in $Z_q(K; \mathbb{G})$ is called a q -cycle and a q -chain in $B_q(K; \mathbb{G})$ is called a q -boundary. For a q -chain A , the $(q-1)$ -chain $\partial(A)$ is also called the boundary of A .

A fundamental fact in homology theory is that $\partial_q \partial_{q+1} = 0$ for any q . This implies that $B_q(K; \mathbb{G}) \subseteq Z_q(K; \mathbb{G})$. The q th homology group of K denoted by $H_q(K; \mathbb{G})$ is defined as the quotient $Z_q(K; \mathbb{G})/B_q(K; \mathbb{G})$. Each coset in $H_q(K; \mathbb{G})$ is called a *homology class* and a cycle is said to be *homologous to* another cycle if they belong to the same homology class. As any boundary cycle represents the homology class 0 in $H_q(K; \mathbb{G})$, a boundary is also said to be *null-homologous*.

The abelian group $\mathbb{G}$ in the above definitions is called the *coefficient group* for the homology groups. Sometimes, when the coefficient group $\mathbb{G}$ is clear, we simply drop it and denote a chain group as $C_q(K)$. This applies to other groups defined in simplicial homology. In this paper, two coefficient groups $\mathbb{Z}_2$ and $\mathbb{Z}$ are used for simplicial homology. When not explicitly stated, the coefficients are assumed to be in $\mathbb{Z}_2$. With $\mathbb{Z}_2$ coefficients, the orientations of simplices no longer matter and a q -chain can be interpreted as a set of q -simplices with summation of two q -chains being the *symmetric difference*. A q -cycle is then a set of q -simplices where every $(q-1)$ -face of these simplices adjoins an even number of q -simplices. Also note that because $\mathbb{Z}_2$ is a field, all groups defined in simplicial homology with $\mathbb{Z}_2$ coefficients become vector spaces and homomorphisms between these groups (such as ∂) become linear maps.

Definition 2.2 (q -weighted). A simplicial complex K is q -weighted if each q -simplex σ of K has a non-negative finite weight $w(\sigma)$. The weight of a q -chain A of K is then defined as $w(A) = \sum_{\sigma \in A} w(\sigma)$.

Definition 2.3 (q -connected). Let K be a simplicial complex, for $q \geq 1$, two q -simplices σ and σ' of K are q -connected in K if there is a sequence of q -simplices of K , $(\sigma_0, \dots, \sigma_l)$, such that $\sigma_0 = \sigma$, $\sigma_l = \sigma'$, and for all $0 \leq i < l$, σ_i and σ_{i+1} share a $(q-1)$ -face. The property of q -connectedness defines an equivalence relation on q -simplices of K . Each set in the partition induced by the equivalence relation constitutes a q -connected component of K . We say K is q -connected if any two q -simplices of K are q -connected in K .

Remark 2.2. See Figure 2a for an example of 1-connected components and 2-connected components.

Definition 2.4 (q -connected cycle). A q -cycle ζ (with $\mathbb{Z}_2$ coefficients) is q -connected if the complex derived by taking the closure of the simplicial set ζ is q -connected.

Persistent homology. We will provide a brief description of persistent homology. We recommend the book by Edelsbrunner and Harer [15] for a detailed explanation of this topic and the book by Chazal et al. [6] for its underlying Mathematical structure, *persistence module*. Note that persistent homology in this paper is always assumed to be with $\mathbb{Z}_2$ coefficients. The persistence algorithm starts with a filtration $\mathcal{F} : \emptyset = K_0 \subseteq K_1 \subseteq \dots \subseteq K_n = K$ of a simplicial complex K , and for each simplex $\sigma_i^{\mathcal{F}}$, inspects whether $\partial(\sigma_i^{\mathcal{F}})$ is a boundary in K_{i-1} . If $\partial(\sigma_i^{\mathcal{F}})$ is a boundary in K_{i-1} , $\sigma_i^{\mathcal{F}}$ is called *positive*; otherwise, it is called *negative*. The d -chains (or d -cycles) in K_i that are not in K_{i-1} are said to be *born in K_i* or *created by $\sigma_i^{\mathcal{F}}$* . A positive d -simplex creates some d -cycles and a negative d -simplex makes some $(d-1)$ -cycles become boundaries. In the latter case, we also say that the negative d -simplex *kills* or *destroys* those $(d-1)$ -cycles. What is central to the persistence algorithm is a notion called *pairing*: A positive simplex is initially *unpaired* when introduced; when a negative d -simplex $\sigma_j^{\mathcal{F}}$ comes, the algorithm finds a $(d-1)$ -cycle created by an unpaired positive $(d-1)$ -simplex $\sigma_i^{\mathcal{F}}$ which is homologous to $\partial(\sigma_j^{\mathcal{F}})$ and *pair $\sigma_j^{\mathcal{F}}$ with $\sigma_i^{\mathcal{F}}$* . Alongside the pairing, a *finite interval* $[j, i)$ is added to the $(d-1)$ th *persistence diagram*, which is denoted by $D_{d-1}(\mathcal{F})$. After all simplices are processed, some positive simplices may still be unpaired. For each $\sigma_i^{\mathcal{F}}$ of these unpaired simplices, an *infinite interval* $[i, +\infty)$ is added to $D_d(\mathcal{F})$, where d is the dimension of $\sigma_i^{\mathcal{F}}$.

Note that the pairing in the persistence algorithm for a given filtration is unique. Also note that in this paper, we assume a filtration of a complex is given and the persistence intervals start and end with indices of the paired simplices. However, in real-life applications, one is often given a function on a simplicial complex. To produce the persistence intervals, a filtration needs to be derived and the endpoints of the intervals are taken as function values on the paired simplices. In such cases, we can associate a given interval to its simplex pair, take the indices of the paired simplices, and get an interval which can serve as an input to our algorithms.

The persistent cycle problems. We can now formally define the minimal persistent cycle problems:

Problem 2.1 (PCYC-FIN $_d$). Given a finite d -weighted simplicial complex K , a filtration $\mathcal{F} : \emptyset = K_0 \subseteq K_1 \subseteq \dots \subseteq K_n = K$, and a finite interval $[\beta, \delta) \in D_d(\mathcal{F})$, this problem asks for computing a d -cycle with the minimal weight which is born in K_β and becomes a boundary in K_δ .

Problem 2.2 (PCYC-INF $_d$). Given a finite d -weighted simplicial complex K , a filtration $\mathcal{F} : \emptyset = K_0 \subseteq K_1 \subseteq \dots \subseteq K_n = K$, and an infinite interval $[\beta, +\infty) \in D_d(\mathcal{F})$, this problem asks for computing a d -cycle with the minimal weight which is born in K_β .

Remark 2.3. The definitions of the above two problems are derived directly from the definition of persistent d -cycles [13].

Undirected flow network. An *undirected flow network* (G, s_1, s_2) consists of an undirected graph G with vertex set $V(G)$ and edge set $E(G)$, a capacity function $c : E(G) \rightarrow [0, +\infty]$, and two non-empty disjoint subsets s_1 and s_2 of $V(G)$. Vertices in s_1 are referred to as *sources* and vertices in s_2 are referred to as *sinks*. A *cut* (S, T) of (G, s_1, s_2) consists of two disjoint subsets S and T of $V(G)$ such that $S \cup T = V(G)$, $s_1 \subseteq S$, and $s_2 \subseteq T$. The set of edges that connect a vertex in S and a vertex in T are referred as the edges *across* the cut (S, T) and is denoted as $\xi(S, T)$. The *capacity* of a cut (S, T) is defined as $c(S, T) = \sum_{e \in \xi(S, T)} c(e)$. A *minimal cut* of (G, s_1, s_2) is a cut with the minimal capacity. Note that we allow parallel edges in G (see Figure 2a) to ease the presentation. These parallel edges can be merged into one edge during computation.




Figure 1: An example of the constructions in our algorithm showing the duality between persistent cycles and cuts having finite capacity for $d = 1$. (a) The input weak 2-pseudomanifold K with its dual flow network drawn in blue, where the central hollow vertex denotes the dummy vertex, the red vertex denotes the source, and all the orange vertices (including the dummy one) denote the sinks. All “dangled” graph edges dual to the outer boundary 1-simplices actually connect to the dummy vertex and these connections are not drawn. (b) The partial complex K_β in the input filtration $\mathcal{F}$, where the bold green 1-simplex denotes $\sigma_\beta^\mathcal{F}$ which creates the green 1-cycle. (c) The partial complex K_δ in $\mathcal{F}$, where the 2-simplex $\sigma_\delta^\mathcal{F}$ creates the pink 2-chain killing the green 1-cycle. (d) The green persistent 1-cycle of the interval $[\beta, \delta)$ is dual to a cut (S, T) having finite capacity, where S contains all the vertices inside the pink 2-chain and T contains all the other vertices. The red graph edges denote those edges across (S, T) and their dual 1-chain is the green persistent 1-cycle.

3 Minimal persistent d -cycles of finite intervals for weak $(d + 1)$ -pseudomanifolds

In this section, we present an algorithm which computes minimal persistent d -cycles for finite intervals given a filtration of a weak $(d + 1)$ -pseudomanifold when $d \geq 1$. The general process is as follows: Suppose that the input weak $(d + 1)$ -pseudomanifold is K associated with a filtration $\mathcal{F} : K_0 \subseteq K_1 \subseteq \dots \subseteq K_n$ and the task is to compute the minimal persistent cycle of a finite interval $[\beta, \delta) \in \mathcal{D}_d(\mathcal{F})$. We first construct an undirected dual graph G for K where vertices of G are dual to $(d + 1)$ -simplices of K and edges of G are dual to d -simplices of K . One dummy vertex termed as *infinite vertex* which does not correspond to any $(d + 1)$ -simplices is added to G for graph edges dual to those boundary d -simplices. We then build an undirected flow network on top of G where the source is the vertex dual to $\sigma_\delta^\mathcal{F}$ and the sink is the infinite vertex along with the set of vertices dual to those $(d + 1)$ -simplices which are added to $\mathcal{F}$ after $\sigma_\delta^\mathcal{F}$. If a d -simplex is $\sigma_\beta^\mathcal{F}$ or added to $\mathcal{F}$ before $\sigma_\beta^\mathcal{F}$, we let the capacity of its dual graph edge be its weight; otherwise, we let the capacity of its dual graph edge be $+\infty$. Finally, we calculate a minimal cut of this flow network and return the d -chain dual to the edges across the minimal cut as a minimal persistent cycle of the interval.

The intuition of the above algorithm is best explained by an example in Figure 1, where $d = 1$. The key to the algorithm is the duality between persistent cycles of the input interval and cuts of the dual flow network having finite capacity. To see this duality, first consider a persistent d -cycle ζ of the input interval $[\beta, \delta)$. There exists a $(d + 1)$ -chain A in K_δ created by $\sigma_\delta^\mathcal{F}$ whose boundary equals ζ , making ζ killed. We can let S be the set of graph vertices dual to the simplices in A and let T be the set of the remaining graph vertices, then (S, T) is a cut. Furthermore, (S, T) must have finite capacity as the edges across it are exactly dual to the d -simplices in ζ and the d -simplices in ζ have indices in $\mathcal{F}$ less than or equal to β . On the other hand, let (S, T) be a cut with finite capacity, then the $(d + 1)$ -chain whose simplices are dual to the vertices in S is created by $\sigma_\delta^\mathcal{F}$. Taking the

boundary of this $(d+1)$ -chain, we get a d -cycle ζ . Because d -simplices of ζ are exactly dual to the edges across (S, T) and each edge across (S, T) has finite capacity, ζ must reside in K_β . We only need to ensure that ζ contains $\sigma_\beta^\mathcal{F}$ in order to show that ζ is a persistent cycle of $[\beta, \delta]$. In Section 3.2, we argue that ζ actually contains $\sigma_\beta^\mathcal{F}$, so ζ is indeed a persistent cycle. Note that while the above explanation introduces the general idea, the rigorous statement and proof of the duality are articulated by Proposition 3.2 and 3.3.

Algorithm 3.1 Computing minimal persistent d -cycles of finite intervals for weak $(d+1)$ -pseudomanifolds

Input:

K : finite d -weighted weak $(d+1)$ -pseudomanifold

d : integer ≥ 1

$\mathcal{F}$: filtration $K_0 \subseteq K_1 \subseteq \dots \subseteq K_n$ of K

$[\beta, \delta]$: finite interval of $D_d(\mathcal{F})$

Output:

minimal persistent d -cycle of $[\beta, \delta]$

```

1: procedure MINPERSCYCFIN( $K, d, \mathcal{F}, [\beta, \delta]$ )
    ▷ set up the complex  $\tilde{K}$  being worked on
2:    $C^{d+1} \leftarrow (d+1)$ -connected component of  $K$  containing  $\sigma_\delta^\mathcal{F}$ 
3:    $\tilde{K} \leftarrow$  closure of the simplicial set  $C^{d+1}$ 
    ▷ construct dual graph
4:    $(G, \theta) \leftarrow \text{DUALGRAPHFIN}(\tilde{K}, d)$ 
    ▷ assign capacity to  $G$ 
5:   for each  $e \in E(G)$  do
6:     if  $\text{ind}(\theta^{-1}(e)) \leq \beta$  then
7:        $c(e) \leftarrow w(\theta^{-1}(e))$ 
8:     else
9:        $c(e) \leftarrow +\infty$ 
    ▷ set the source
10:   $s_1 \leftarrow \{\theta(\sigma_\delta^\mathcal{F})\}$ 
    ▷ set the sink
11:   $s_2 \leftarrow \{v \in V(G) \mid v \neq \phi, \text{ind}(\theta^{-1}(v)) > \delta\}$ 
12:  if  $\phi \in V(G)$  then
13:     $s_2 \leftarrow s_2 \cup \{\phi\}$ 
14:   $(S^*, T^*) \leftarrow$  min-cut of  $(G, s_1, s_2)$ 
15:  return  $\theta^{-1}(\xi(S^*, T^*))$ 

```

We list the pseudo-code in Algorithm 3.1 and it works as follows: Line 2 and 3 set up a complex $\tilde{K}$ that the algorithm mainly works on, where $\tilde{K}$ is taken as the closure of the $(d+1)$ -connected component of K containing $\sigma_\delta^\mathcal{F}$. The reason for working on $\tilde{K}$ instead of the entire complex is explained later in this section. Line 4 constructs the dual graph G from $\tilde{K}$ and line 5–13 builds the flow network on top of G . Note that we denote the infinite vertex by ϕ . Line 14 computes a minimal cut for the flow network and line 15 returns the d -chain dual to the edges across the minimal cut. In the pseudo-codes of this paper, to ease the exposition, we treat a Mathematical

function as a computer program object. For example, the function θ returned by DUALGRAPHFIN in Algorithm 3.1 denotes the bijection between the simplices of $\tilde{K}$ and their dual vertices or edges (see Section 3.1 for details). In practice, these constructs can be easily implemented in any computer programming language.

To see the reason why we work on $\tilde{K}$, we first note that the dual graph constructed directly from K may be disconnected[‡]. While cuts are still well-defined for a disconnected flow network, one may prefer a connected one as the minimal cut computation only concerns the graph component containing the source. By constructing the dual graph from $\tilde{K}$, it can be ensured that the graph is connected. In order for Algorithm 3.1 to work, one has to further show that the sink is non-empty so that the computed persistent cycle is non-empty. This is verified in Proposition 3.1. An intuitive reason why the computation from $\tilde{K}$ is still correct is as follows: Each persistent d -cycle ζ of the given interval corresponds to a $(d+1)$ -chain A which kills ζ , i.e., $\partial(A) = \zeta$. Suppose that A is not entirely contained in $\tilde{K}$. Notice that $A \cap \tilde{K} \neq \emptyset$ and contains at least the killer simplex $\sigma_\delta^{\mathcal{F}}$. Then $\partial(A \cap \tilde{K})$ must be a persistent cycle of the interval residing in $\tilde{K}$ which has a smaller weight. Hence, a minimal persistent cycle must reside in $\tilde{K}$. In Section 3.2, we formally verify the construction.

Complexity. The time complexity of Algorithm 3.1 depends on the encoding scheme of the input and the data structure used for representing a simplicial complex. For encodings of the input, we assume K and $\mathcal{F}$ to be represented by a sequence of all the simplices of K ordered by their indices in $\mathcal{F}$, where each simplex is denoted by its set of vertices. We also assume a simple yet reasonable simplicial complex data structure as follows: In each dimension, simplices are mapped to integral identifiers ranging from 0 to the number of simplices in that dimension minus 1; each q -simplex has an array (or linked list) storing all the id's of its $(q+1)$ -cofaces; a hash map for each dimension is maintained for the query of the integral id of each simplex in that dimension based on the spanning vertices of the simplex. We further assume d to be constant. By the above assumptions, let n be the size (number of bits) of the encoded input, then there are no more than n elementary $O(1)$ operations in line 2 and 3. So, the time complexity of line 2 and 3 is $O(n)$. It is not hard to verify that the flow network construction also takes $O(n)$ time so the time complexity of Algorithm 3.1 is determined by the minimal cut algorithm. Using the max-flow algorithm by Orlin [25], the time complexity of Algorithm 3.1 becomes $O(n^2)$.

In the rest of this section, we first explain the bijection θ returned by DUALGRAPHFIN, then prove the correctness of the algorithm.

3.1 The bijection θ

The vertex set $V(G)$ of G contains vertices which correspond to the $(d+1)$ -simplices of $\tilde{K}$. The set $V(G)$ may also contain an infinite vertex ϕ if $\tilde{K}$ contains any boundary d -simplex. We define a bijection

$$\theta : \{(d+1)\text{-simplices of } \tilde{K}\} \rightarrow V(G) \setminus \{\phi\}$$

such that for any $(d+1)$ -simplex σ^{d+1} of $\tilde{K}$, $\theta(\sigma^{d+1})$ is the vertex that σ^{d+1} is dual to. Similarly, we define another bijection

$$\theta : \{d\text{-simplices of } \tilde{K}\} \rightarrow E(G)$$

using the same notation θ .

[‡]For an example in $d = 1$, take K as two disconnected triangulated 2-spheres. Its dual graph consists of two connected components.

Note that we can take the image of a subset of the domain under a function. Therefore, if (S, T) is a cut for a flow network built on G , then $\theta^{-1}(\xi(S, T))$ denotes the set of d -simplices dual to the edges across the cut. Also note that since simplicial chains with $\mathbb{Z}_2$ coefficients can be interpreted as sets, $\theta^{-1}(\xi(S, T))$ is also a d -chain.

3.2 Algorithm correctness

In this subsection, we prove the correctness of Algorithm 3.1. Some of the symbols we use refer to Algorithm 3.1.

Proposition 3.1. *In Algorithm 3.1, the sink s_2 is not an empty set.*

Proof. For contradiction, suppose that s_2 is an empty set. Then, $\phi \notin V(G)$ and $\sigma_\delta^{\mathcal{F}}$ is the $(d+1)$ -simplex of $\tilde{K}$ with the greatest index in $\mathcal{F}$. Because $\phi \notin V(G)$, any d -simplex of $\tilde{K}$ must be a face of two $(d+1)$ -simplices of $\tilde{K}$, so the set of $(d+1)$ -simplices of $\tilde{K}$ forms a $(d+1)$ -cycle created by $\sigma_\delta^{\mathcal{F}}$. Then $\sigma_\delta^{\mathcal{F}}$ must be a positive simplex in $\mathcal{F}$, which is a contradiction. $\square$

The following two propositions specify the duality mentioned at the beginning of this section:

Proposition 3.2. *For any cut (S, T) of (G, s_1, s_2) with finite capacity, the d -chain $\zeta = \theta^{-1}(\xi(S, T))$ is a persistent d -cycle of $[\beta, \delta)$ and $w(\zeta) = c(S, T)$.*

Proof. Let $A = \theta^{-1}(S)$, we first want to prove $\zeta = \partial(A)$, so that ζ is a cycle. Let σ^d be any d -simplex of ζ , then $\theta(\sigma^d)$ connects a vertex $u \in S$ and a vertex $v \in T$. If $v = \phi$, then σ^d cannot be a face of another $(d+1)$ -simplex in K other than $\theta^{-1}(u)$. So, σ^d is a face of exactly one $(d+1)$ -simplex of A . If $v \neq \phi$, then σ^d is also a face of exactly one $(d+1)$ -simplex of A . Therefore, $\sigma^d \in \partial(A)$. On the other hand, let σ^d be any d -simplex of $\partial(A)$, then σ^d is a face of exactly one $(d+1)$ -simplex σ_0^{d+1} of A . If σ^d is a face of another $(d+1)$ -simplex σ_1^{d+1} in K , then $\sigma_1^{d+1} \in \tilde{K}$ and $\sigma_1^{d+1} \notin A$. So, $\theta(\sigma^d)$ connects the vertex $\theta(\sigma_0^{d+1}) \in S$ and the vertex $\theta(\sigma_1^{d+1}) \in T$ in the graph G . If σ^d is a face of exactly one $(d+1)$ -simplex in K , $\theta(\sigma^d)$ must connect $\theta(\sigma_0^{d+1}) \in S$ and $\phi \in T$ in G . So we have $\theta(\sigma^d) \in \xi(S, T)$, i.e., $\sigma^d \in \theta^{-1}(\xi(S, T))$.

We then show that ζ is created by $\sigma_\beta^{\mathcal{F}}$. By Proposition 3.1, ζ cannot be empty. Therefore, for contradiction, we can suppose that ζ is created by a d -simplex $\sigma^d \neq \sigma_\beta^{\mathcal{F}}$. Because $c(S, T)$ has finite capacity, we have that $\text{ind}(\sigma^d) < \beta$. We can let ζ' be a persistent cycle of $[\beta, \delta)$ and $\zeta' = \partial(A')$ where A' is a $(d+1)$ -chain of K_δ . Then we have $\zeta + \zeta' = \partial(A + A')$. Since A and A' are both created by $\sigma_\delta^{\mathcal{F}}$, then $A + A'$ is created by a $(d+1)$ -simplex with an index less than δ in $\mathcal{F}$. So $\zeta + \zeta'$ is a d -cycle created by $\sigma_\beta^{\mathcal{F}}$ which becomes a boundary before $\sigma_\delta^{\mathcal{F}}$ is added. This means that $\sigma_\beta^{\mathcal{F}}$ is already paired when $\sigma_\delta^{\mathcal{F}}$ is added, contradicting the fact that $\sigma_\beta^{\mathcal{F}}$ is paired with $\sigma_\delta^{\mathcal{F}}$. Similarly, we can prove that ζ is not a boundary until $\sigma_\delta^{\mathcal{F}}$ is added, so ζ is a persistent cycle of $[\beta, \delta)$. Since (S, T) has finite capacity, we must have

$$c(S, T) = \sum_{e \in \theta(\zeta)} c(e) = \sum_{\theta^{-1}(e) \in \zeta} w(\theta^{-1}(e)) = w(\zeta) \quad \square$$

Proposition 3.3. *For any persistent d -cycle ζ of $[\beta, \delta)$, there exists a cut (S, T) of (G, s_1, s_2) such that $c(S, T) \leq w(\zeta)$.*

Proof. Let A be a $(d+1)$ -chain in K_δ such that $\zeta = \partial(A)$. Note that A is created by $\sigma_\delta^{\mathcal{F}}$ and ζ is the set of d -simplices which are face of exactly one $(d+1)$ -simplex of A . Let $\zeta' = \zeta \cap \tilde{K}$ and $A' = A \cap \tilde{K}$, we claim that $\zeta' = \partial(A')$. To prove this, first let σ^d be any d -simplex of ζ' , then σ^d

is a face of exactly one $(d+1)$ -simplex σ^{d+1} of A . Since $\sigma^d \in \tilde{K}$, it is also true that $\sigma^{d+1} \in \tilde{K}$, so $\sigma^{d+1} \in A'$. Then σ^d is a face of exactly one $(d+1)$ -simplex of A' , so $\sigma^d \in \partial(A')$. On the other hand, let σ^d be any d -simplex of $\partial(A')$, then σ^d is a face of exactly one $(d+1)$ -simplex σ_0^{d+1} of A' . Note that $\sigma_0^{d+1} \in A$ and we then want to prove that σ^d is a face of exactly one $(d+1)$ -simplex σ_0^{d+1} of A . Suppose that σ^d is a face of another $(d+1)$ -simplex σ_1^{d+1} of A , then $\sigma_1^{d+1} \in \tilde{K}$ because $\sigma_0^{d+1} \in \tilde{K}$. So we have $\sigma_1^{d+1} \in A \cap \tilde{K} = A'$, contradicting the fact that σ^d is a face of exactly one $(d+1)$ -simplex of A' . Then we have $\sigma^d \in \partial(A)$. Since $\sigma_0^{d+1} \in \tilde{K}$, we have $\sigma^d \in \tilde{K}$, which means that $\sigma^d \in \zeta'$.

Let $S = \theta(A')$ and $T = V(G) \setminus S$, then it is true that (S, T) is a cut of (G, s_1, s_2) because A' is created by σ_δ^F . We claim that $\theta^{-1}(\xi(S, T)) = \partial(A')$. The proof of the equality is similar to the one in the proof of Proposition 3.2. It follows that $\xi(S, T) = \theta(\zeta')$. We then have that

$$c(S, T) = \sum_{e \in \theta(\zeta')} c(e) = \sum_{\theta^{-1}(e) \in \zeta'} w(\theta^{-1}(e)) = w(\zeta')$$

because each d -simplex of ζ' has an index less than or equal to β in $\mathcal{F}$.

Finally, because ζ' is a subchain of ζ , we must have $c(S, T) = w(\zeta') \leq w(\zeta)$. $\square$

Combining the above facts, we can conclude:

Theorem 3.1. *Algorithm 3.1 computes a minimal persistent d -cycle for the given interval $[\beta, \delta)$.*

Proof. First, the flow network (G, s_1, s_2) constructed by Algorithm 3.1 must be valid by Proposition 3.1. Next, because the interval $[\beta, \delta)$ must have a persistent cycle, by Proposition 3.3, the flow network (G, s_1, s_2) has a cut with finite capacity. This means that $c(S^*, T^*)$ is finite. By Proposition 3.2, the chain $\zeta^* = \theta^{-1}(\xi(S^*, T^*))$ is a persistent cycle of $[\beta, \delta)$. Assume that ζ^* is not a minimal persistent cycle of $[\beta, \delta)$ and instead let ζ' be a minimal persistent cycle of $[\beta, \delta)$. Then there exists a cut (S', T') such that $c(S', T') \leq w(\zeta') < w(\zeta^*) = c(S^*, T^*)$ by Proposition 3.2 and 3.3, contradicting the fact that (S^*, T^*) is a minimal cut. $\square$

4 Minimal persistent d -cycles of infinite intervals for weak $(d+1)$ -pseudomanifolds embedded in $\mathbb{R}^{d+1}$

We already mentioned that computing minimal persistent d -cycles ($d \geq 2$) for infinite intervals is NP-hard even if we restrict to weak $(d+1)$ -pseudomanifolds (see Section 5.3 for a proof). However, when the complex is embedded in $\mathbb{R}^{d+1}$, the problem becomes polynomially tractable. In this section, we present an algorithm for this problem in $d \geq 1$ [§]. The algorithm uses a similar duality described in Section 3. However, a direct use of the approach in Section 3 does not work. For example, in Figure 2a, 1-simplices that do not have any 2-cofaces cannot reside in any 2-connected component of the given complex. Hence, no cut in the flow network may correspond to a persistent cycle of the infinite interval created by such a 1-simplex. Furthermore, unlike the finite interval case, we do not have a negative simplex whose dual can act as a source in the flow network.

Let $(K, \mathcal{F}, [\beta, +\infty))$ be an input to the problem where K is a weak $(d+1)$ -pseudomanifold embedded in $\mathbb{R}^{d+1}$, $\mathcal{F} : K_0 \subseteq K_1 \subseteq \dots \subseteq K_n$ is a filtration of K , and $[\beta, +\infty)$ is an infinite interval of $D_d(\mathcal{F})$. By the definition of the problem, the task boils down to computing a minimal d -cycle containing σ_β^F in K_β . Note that K_β is also a weak $(d+1)$ -pseudomanifold embedded in $\mathbb{R}^{d+1}$.

[§]As mentioned earlier, when $d = 1$, this problem is polynomially tractable for arbitrary complexes.




Figure 2: (a) A weak 2-pseudomanifold $\tilde{K}$ embedded in $\mathbb{R}^2$ with three voids. Its dual graph is drawn in blue. The complex has one 1-connected component and four 2-connected components with the 2-simplices in different 2-connected components colored differently. (b) An example illustrating the pairing of boundary d -simplices in the neighborhood of a $(d-1)$ -simplex for $d=1$. The four boundary 1-simplices produce six oriented boundary 1-simplices and the paired oriented 1-simplices are colored the same.

Generically, assume $\tilde{K}$ is an arbitrary weak $(d+1)$ -pseudomanifold embedded in $\mathbb{R}^{d+1}$ and we want to compute a minimal d -cycle containing a d -simplex $\tilde{\sigma}$ for $\tilde{K}$. By the embedding assumption, the connected components of $\mathbb{R}^{d+1} \setminus |\tilde{K}|$ are well defined and we call them the *voids* of $\mathbb{R}^{d+1} \setminus |\tilde{K}|$. The complex $\tilde{K}$ has a natural (undirected) dual graph structure as exemplified by Figure 2a for $d=1$, where the graph vertices are dual to the $(d+1)$ -simplices as well as the voids and the graph edges are dual to the d -simplices. The duality between cycles and cuts is as follows: Since the ambient space $\mathbb{R}^{d+1}$ is contractible (homotopy equivalent to a point), every d -cycle in $\tilde{K}$ is the boundary of a $(d+1)$ -dimensional region obtained by point-wise union of certain $(d+1)$ -simplices and/or voids. We can derive a cut[¶] of the dual graph by putting all vertices contained in the $(d+1)$ -dimensional region into one vertex set and putting the rest into the other vertex set. On the other hand, for every cut of the graph, we can take the point-wise union of all the $(d+1)$ -simplices and voids dual to the graph vertices in one set of the cut and derive a $(d+1)$ -dimensional region. The boundary of the derived $(d+1)$ -dimensional region is then a d -cycle in $\tilde{K}$. We observe that by making the source and sink dual to the two $(d+1)$ -simplices or voids that $\tilde{\sigma}$ adjoins, we can build a flow network where a minimal cut produces a minimal d -cycle in $\tilde{K}$ containing $\tilde{\sigma}$.

The efficiency of the above algorithm is in part determined by the efficiency of the dual graph construction. This step requires identifying the voids that the boundary d -simplices are incident on. A straightforward approach would be to first group the boundary d -simplices into d -cycles by local geometry, and then build the nesting structure of these d -cycles to correctly reconstruct the boundaries of the voids. This approach has a quadratic worst-case complexity. To make the void boundary reconstruction faster, we assume that the simplicial complex being worked on is d -connected so that building the nesting structure is not needed. Our reconstruction then runs in almost linear time. To satisfy the d -connected assumption, we begin our algorithm by taking $\tilde{K}$ as a d -connected subcomplex of K_β containing σ_β^F and continue only with this $\tilde{K}$. The computed output is still correct because the minimal cycle in $\tilde{K}$ is again a minimal cycle in K_β as shown in Section 4.2.

We list the pseudo-code in Algorithm 4.1 and it works as follows: Line 2–5 set up the complex

[¶]The cut here is defined on a graph without sources and sinks, so the cut is simply a partition of the vertex set into two sets.

Algorithm 4.1 Computing minimal persistent d -cycles of infinite intervals for weak $(d+1)$ -pseudomanifolds embedded in $\mathbb{R}^{d+1}$

Input:

K : finite d -weighted weak $(d+1)$ -pseudomanifold embedded in $\mathbb{R}^{d+1}$
 d : integer ≥ 1
 $\mathcal{F}$: filtration $K_0 \subseteq K_1 \subseteq \dots \subseteq K_n$ of K
 $[\beta, +\infty)$: infinite interval of $D_d(\mathcal{F})$

Output:

minimal persistent d -cycle of $[\beta, +\infty)$

```

1: procedure MINPERSCYCINF( $K, d, \mathcal{F}, [\beta, +\infty)$ )
    ▷ set up the complex  $\tilde{K}$  being worked on
2:    $K'_\beta \leftarrow \text{PRUNE}(K_\beta, d)$ 
3:    $C_\beta \leftarrow d$ -connected component of  $K'_\beta$  containing  $\sigma_\beta^\mathcal{F}$ 
4:    $\Sigma^{d+1} \leftarrow \{\sigma \in K'_\beta \mid \sigma \text{ is a } (d+1)\text{-simplex and all } d\text{-faces of } \sigma \text{ are in } C_\beta\}$ 
5:    $\tilde{K} \leftarrow (\text{closure of the simplicial set } C_\beta) \cup \Sigma^{d+1}$ 
    ▷ construct dual graph
6:    $(\vec{\zeta}_1, \dots, \vec{\zeta}_k) \leftarrow \text{VOIDBOUNDARY}(\tilde{K}, d)$ 
7:    $(G, \theta) \leftarrow \text{DUALGRAPHINF}(\tilde{K}, d, \vec{\zeta}_1, \dots, \vec{\zeta}_k)$ 
    ▷ assign capacity to  $G$ 
8:   for each  $e \in E(G)$  do
9:      $c(e) \leftarrow w(\theta^{-1}(e))$ 
10:   $(v_1, v_2) \leftarrow$  end vertices of edge  $\theta(\sigma_\beta^\mathcal{F})$  in  $G$ 
    ▷ set the source
11:   $s_1 \leftarrow \{v_1\}$ 
    ▷ set the sink
12:   $s_2 \leftarrow \{v_2\}$ 
13:   $(S^*, T^*) \leftarrow$  min-cut of  $(G, s_1, s_2)$ 
14:  return  $\theta^{-1}(\xi(S^*, T^*))$ 

```

$\tilde{K}$ that the algorithm works on. Line 2 prunes K_β to produce a complex K'_β . Given (K_β, d) , the PRUNE subroutine iteratively deletes a d -simplex σ^d of K_β such that there is a $(d-1)$ -face of σ^d having σ^d as the only d -coface (i.e., σ^d is a dangled d -simplex), until no such d -simplex can be found. It is not hard to verify that PRUNE only deletes d -simplices not residing in any d -cycles, so a minimal d -cycle containing $\sigma_\beta^\mathcal{F}$ is never deleted. We perform the pruning because it can reduce the graph size for the minimal cut computation which is more time consuming. In line 3–5, we take the d -connected component C_β of K'_β containing $\sigma_\beta^\mathcal{F}$ and add a set Σ^{d+1} of $(d+1)$ -simplices to the closure of C_β to form $\tilde{K}$. The set Σ^{d+1} contains all $(d+1)$ -simplices of K'_β whose d -faces reside in C_β . The reason of adding the set Σ^{d+1} is to reduce the number of voids for the complex $\tilde{K}$ and in turn reduce the running time of the subsequent void boundary reconstruction. For example, in Figure 3b, we could treat the entire complex as K'_β , all 1-simplices as C_β , and all 2-simplices as Σ^{d+1} . If we do not add Σ^{d+1} to the closure of C_β , there will be seven more voids corresponding to the seven 2-simplices. Line 6 reconstructs the void boundaries for $\tilde{K}$. Each returned $\vec{\zeta}_j$ denotes a set of d -simplices forming the boundary of a void. As indicated in Section 4.1, the d -simplices




Figure 3: Examples showing how the void boundaries are reconstructed for $d = 1$. (a) Oriented boundary 1-simplices (drawn as dashed edges) of a simplicial complex are grouped into six 1-cycles and these six 1-cycles are further grouped into four void boundaries with each void boundary identically colored. (b) With the complex being 1-connected, the four grouped 1-cycles are exactly the boundaries of the four voids.

in a void boundary are oriented. Line 7 constructs the dual graph G based on the reconstructed void boundaries. Similar to Algorithm 3.1, the function θ returned by `DUALGRAPHINF` denotes the bijection from d -simplices of $\tilde{K}$ to $E(G)$. Line 8–12 build the flow network on top of G . The capacity of each edge is equal to the weight of its dual d -simplex and the source and sink are selected as previously described. Line 13 computes a minimal cut for the flow network and line 14 returns the d -chain dual to the edges across the minimal cut.

Complexity. We make the same assumptions as in the complexity analysis for Algorithm 3.1. Since the void boundary reconstruction needs to sort the d -cofaces of certain $(d - 1)$ -simplices, its worst-case time complexity is $O(n \log n)$. Then, all operations other than the minimal cut computation take $O(n \log n)$ time. Therefore, similar to Algorithm 3.1, Algorithm 4.1 achieves a complexity of $O(n^2)$ by using Orlin’s max-flow algorithm [25].

In the rest of this section, we first describe the subroutine `VOIDBOUNDARY` invoked by Algorithm 4.1 and then prove the correctness of the algorithm.

4.1 Void boundary reconstruction

As previously stated, the object of the reconstruction is to identify which voids a boundary d -simplex of $\tilde{K}$ is incident on. The task becomes complicated because a void may have disconnected boundaries and a d -simplex may bound more than one void. This is exemplified in Figure 3a. To address this issue, we orient the boundary d -simplices and determine the orientations consistently from the voids they bound. This is possible because an orientation of a d -simplex in $\mathbb{R}^{d+1}$ associates exactly one of its two sides to the d -simplex. To reconstruct the boundaries, we first inspect the neighborhood of each $(d - 1)$ -simplex being a face of a boundary d -simplex and pair the oriented boundary d -simplices in the neighborhood which locally bound the same void. Figure 2b gives an example of the oriented boundary d -simplices pairing for $d = 1$. In Figure 2b, there are three local voids each colored differently. The oriented 1-simplices with the same color bound the same void and are paired.

After pairing the oriented boundary d -simplices, we group them by putting paired ones into the same group. Each group then forms a d -cycle (with $\mathbb{Z}$ coefficients). This is exemplified by Figure 3

for $d = 1$. Note that in general, the above grouping does not fully reconstruct the void boundaries. This can be seen from Figure 3a where the complex has four voids but the grouping produces six 1-cycles. In order to fully reconstruct the boundaries, one has to retrieve the nesting structure of these d -cycles, which may take $\Omega(n^2)$ time in the worst-case. However, as we work on a complex $\tilde{K}$ that is d -connected, we cannot have voids with disconnected boundaries. Therefore, the grouping of oriented d -simplices can fully recover the void boundaries. Figure 3b gives an example for this when $d = 1$, where we add two 1-simplices to make the complex 1-connected. The four 1-cycles produced by the grouping are exactly the boundaries of the four voids.

In the rest of this subsection, we formalize the above ideas for reconstructing void boundaries and provide a proof for the correctness. Throughout this subsection, $\tilde{K}$ and d are as defined in Algorithm 4.1. We first introduce the definition of the natural orientation of a q -simplex in $\mathbb{R}^q$. We use its induced orientation to canonically orient the boundary simplices.

Definition 4.1 (Natural orientation [22]). Let $q > 1$ and $\sigma = \{v_0, \dots, v_q\}$ be a q -simplex in $\mathbb{R}^q$, an oriented simplex $\vec{\sigma} = [v'_0, \dots, v'_q]$ of σ is *naturally oriented* if $\det(v'_1 - v'_0, \dots, v'_q - v'_0) > 0$. For each face σ' of σ , the natural orientation of σ induces an orientation of σ' which we term as the *induced orientation*.

We now formally define the boundary of a void as follows:

Definition 4.2 (Boundary of void). Let K be a simplicial complex embedded in $\mathbb{R}^q$ where $q \geq 2$, an oriented $(q-1)$ -simplex $\vec{\sigma}^{q-1} = [v_0, \dots, v_{q-1}]$ of K is said to *bound* a void $\mathcal{V}$ of $\mathbb{R}^q \setminus |K|$ if the following conditions are satisfied:

- The simplex $\sigma^{q-1} = \{v_0, \dots, v_{q-1}\}$ is contained in the closure of $\mathcal{V}$.
- Let u be an interior point of $\sigma^{q-1} = \{v_0, \dots, v_{q-1}\}$, v be a point in $\mathcal{V}$ such that the line segment $\overline{uv}$ is contained in $\mathcal{V}$ and $\overline{uv}$ is orthogonal to the hyperplane spanned by σ^{q-1} . Furthermore, let $\vec{\sigma}^q$ be the naturally oriented simplex of $\{v, v_0, \dots, v_{q-1}\}$. Then, $\vec{\sigma}^{q-1}$ has the induced orientation from $\vec{\sigma}^q$.

The *boundary* of a void $\mathcal{V}$ is then defined as the set of oriented $(q-1)$ -simplices of K bounding $\mathcal{V}$.

Remark 4.1. We can also interpret the boundary of a void as a sum of oriented $(q-1)$ -simplices, then the boundary defines a $(q-1)$ -cycle (with $\mathbb{Z}$ coefficients).

We now describe the pairing algorithm of the oriented boundary d -simplices for $\tilde{K}$. From now on, we denote the set of boundary d -simplices of $\tilde{K}$ as $\text{bd}(\tilde{K})$. Let σ^{d-1} be a $(d-1)$ -simplex which is a face of a d -simplex in $\text{bd}(\tilde{K})$, we first take a 2D plane Δ which contains an interior point of σ^{d-1} and is orthogonal to the hyperplane spanned by σ^{d-1} . We then take the intersection of the plane Δ with each boundary d -simplex in the neighborhood of σ^{d-1} to get a set of line segments that we order circularly starting from an arbitrary one. For each two consecutive line segments in this order which enclose a void, we pick a point p on the plane Δ which resides in the void. Suppose that one of the two line segments is derived from a boundary d -simplex $\sigma_0^d = \{v_0, \dots, v_d\}$. We take the $(d+1)$ -simplex $\sigma^{d+1} = \{p, v_0, \dots, v_d\}$ and the induced oriented simplex $\vec{\sigma}_0^d$ of σ_0^d derived from the naturally oriented simplex of σ^{d+1} . For the other line segment, we similarly derive an induced oriented simplex $\vec{\sigma}_1^d$ and pair the two oriented d -simplices $\vec{\sigma}_0^d$ and $\vec{\sigma}_1^d$. Figure 2b can be reused to exemplify the pairing. The union of the shaded regions in the figure is the plane Δ and a, b, c , and d are the line segments derived from intersecting the plane with four boundary d -simplices. Taking the circular order a, b, c, d , we see that the consecutive ones which enclose a void are (a, b) , (c, d) ,

and (d, a) . For (a, b) , we can pick p as an interior point in the blue region and the two oriented d -simplices corresponding to a and b can be induced and paired.

In summary, the steps of the VOIDBOUNDARY subroutine are the following:

1. For each $(d - 1)$ -simplex σ^{d-1} being a face of a d -simplex in $\text{bd}(\tilde{K})$, pair all oriented boundary d -simplices in the neighborhood.
2. After gathering all the pairing, group the oriented boundary d -simplices by putting all paired ones into a group.
3. Return $(\vec{\zeta}_1, \dots, \vec{\zeta}_k)$, each of which is a group of the oriented boundary d -simplices.

The following theorem concludes the correctness of the reconstruction:

Theorem 4.1. *Any $\vec{\zeta}_j$ returned by VOIDBOUNDARY is the boundary of a void of $\mathbb{R}^{d+1} \setminus |\tilde{K}|$.*

Proof. See Appendix A. □

4.2 Algorithm correctness

To prove the correctness of Algorithm 4.1, we need two conclusions about cycles with $\mathbb{Z}_2$ coefficients. Specifically, Proposition 4.1 says that an embedded $(q - 1)$ -cycle in $\mathbb{R}^q$ separates the space and hence the two oriented simplices of a $(q - 1)$ -simplex in the cycle bound different voids. Proposition 4.2 says that a q -simplex in a q -cycle belongs to a q -connected sub-cycle of the q -cycle.

Proposition 4.1. *Let $q \geq 2$, ζ be a $(q - 1)$ -cycle (with $\mathbb{Z}_2$ coefficients) of a simplicial complex embedded in $\mathbb{R}^q$, and $\mathcal{Z}$ be the closure of the simplicial set ζ . Then for any $(q - 1)$ -simplex σ of ζ , the two oriented simplices of σ must bound different voids of $\mathbb{R}^q \setminus |\mathcal{Z}|$.*

Proof. Consider a closed topological q -ball $\mathbb{B}$ such that $\sigma \subseteq \mathbb{B}$ and $\mathbb{B} \cap |\mathcal{Z} \setminus \sigma|$ equals the boundary of σ . Let $\mathbb{B}_1$ and $\mathbb{B}_2$ be the two open half balls of $\mathbb{B}$ separated by σ . Then it is true that the two oriented simplices of σ bound different voids of $\mathbb{R}^q \setminus |\mathcal{Z}|$ if and only if $\mathbb{B}_1$ and $\mathbb{B}_2$ are not connected in $\mathbb{R}^q \setminus |\mathcal{Z}|$. So we only need to show that $\mathbb{B}_1$ and $\mathbb{B}_2$ are not connected in $\mathbb{R}^q \setminus |\mathcal{Z}|$. Consider a filtration of $\mathcal{Z}$ where σ is the last simplex added. Because σ is a positive simplex in the filtration, by adding σ , the dimension of H_{q-1} must increase by 1. By Alexander duality, the dimension of H_0 of the complement space also increases by 1. Then $\mathbb{B}_1$ and $\mathbb{B}_2$ cannot be connected in $\mathbb{R}^q \setminus |\mathcal{Z}|$. □

Proposition 4.2. *Let ζ be a q -cycle (with $\mathbb{Z}_2$ coefficients) of a simplicial complex where $q > 0$, then for any q -simplex σ of ζ , there must be a q -cycle ζ' (with $\mathbb{Z}_2$ coefficients) containing σ such that $\zeta' \subseteq \zeta$ and ζ' is q -connected.*

Proof. We can construct an undirected graph L for ζ , with vertices of L corresponding to the q -simplices in ζ . For each $(q - 1)$ -simplex σ^{q-1} which is a face of a q -simplex of ζ , let $\mathcal{N}$ be the set of q -simplices in ζ having σ^{q-1} as a face, then $|\mathcal{N}|$ must be even. We can pair q -simplices of $\mathcal{N}$ arbitrarily, and make each pair of q -simplices form an edge in L . Let C be the connected component of L containing the corresponding vertex of σ and ζ' be the q -chain corresponding to C , then ζ' must be a cycle. This is because we can pair the $(q - 1)$ -faces of all q -simplices in ζ' according to the edges in L , so $\partial(\zeta') = 0$. Furthermore, ζ' contains σ , $\zeta' \subseteq \zeta$, and ζ' is q -connected. □

Throughout the rest of this subsection, some of the symbols we use refer to Algorithm 4.1. We endow the ambient space $\mathbb{R}^{d+1}$ with a “cellular complex” structure by treating voids of $\mathbb{R}^{d+1} \setminus |\tilde{K}|$ as $(d + 1)$ -dimensional “cells”. This cellular complex of $\mathbb{R}^{d+1}$ is denoted as $\mathcal{R}^{d+1}$ and $\mathcal{R}^{d+1} =$

$\tilde{K} \cup \{\text{voids of } \mathbb{R}^{d+1} \setminus |\tilde{K}|\}$. For $\mathcal{R}^{d+1}$, most terminologies from algebraic topology for simplicial complexes are inherited with the exception that $(d+1)$ -dimensional elements of $\mathcal{R}^{d+1}$ are called $(d+1)$ -cells. Then, we can also let θ denote the bijection from $(d+1)$ -cells of $\mathcal{R}^{d+1}$ to $V(G)$. To derive $\partial(\mathcal{V})$ for a void $\mathcal{V}$ of $\mathbb{R}^{d+1} \setminus |\tilde{K}|$, we map oriented d -simplices in the boundary of $\mathcal{V}$ (Definition 4.2) to their corresponding unoriented d -simplices. Then $\partial(\mathcal{V})$ is defined as the sum (with $\mathbb{Z}_2$ coefficients) of these unoriented d -simplices. It is not hard to see that $\partial(\mathcal{V})$ is a d -cycle (with $\mathbb{Z}_2$ coefficients) because each void boundary is a d -cycle (with $\mathbb{Z}$ coefficients).

Proposition 4.3. *For any cut (S, T) of (G, s_1, s_2) , the d -chain $\zeta = \theta^{-1}(\xi(S, T))$ is a persistent d -cycle of $[\beta, +\infty)$ and $w(\zeta) = c(S, T)$.*

Proof. We have three things to show: (i) ζ contains $\sigma_\beta^{\mathcal{F}}$; (ii) $w(\zeta) = c(S, T)$; (iii) ζ is a cycle. Claim (i) and (ii) are not hard to verify and we prove claim (iii) by showing that $\zeta = \sum_{\alpha \in \theta^{-1}(S)} \partial(\alpha)$, so that as a sum of cycles, ζ is a cycle. The detail for the equality of the two chains is omitted as it is similar to the one in the proof of Proposition 3.2. $\square$

Proposition 4.4. *For any persistent d -cycle ζ of $[\beta, +\infty)$, there exists a cut (S, T) of (G, s_1, s_2) such that $c(S, T) \leq w(\zeta)$.*

Proof. Because of the nature of the pruning, ζ must reside in K'_β . By Proposition 4.2, there must be a d -cycle $\zeta' \subseteq \zeta$ such that ζ' is d -connected and contains $\sigma_\beta^{\mathcal{F}}$. Hence, ζ' resides in $\tilde{K}$. Let $\mathcal{Z}'$ be the closure of the simplicial set ζ' , we can run the void boundary reconstruction algorithm of Section 4.1 on $\mathcal{Z}'$ and take a void boundary $\vec{\zeta}$ containing an oriented simplex $\vec{\sigma}_\beta^{\mathcal{F}}$ of $\sigma_\beta^{\mathcal{F}}$. We can map each oriented simplex of $\vec{\zeta}$ to its unoriented simplex and let ζ_0 be the sum of these unoriented simplices, then ζ_0 is a d -cycle (with $\mathbb{Z}_2$ coefficients) and $\zeta_0 \subseteq \zeta'$. By Proposition 4.1, the oppositely oriented simplex of $\vec{\sigma}_\beta^{\mathcal{F}}$ must not be in $\vec{\zeta}$, so ζ_0 contains $\sigma_\beta^{\mathcal{F}}$. Let $\vec{\zeta}$ bound a void $\mathcal{V}$ of $\mathbb{R}^{d+1} \setminus |\mathcal{Z}'|$, we can let $\mathcal{A}$ be the $(d+1)$ -chain of $\mathcal{R}^{d+1}$ consisting of all the $(d+1)$ -cells residing in $\mathcal{V}$ and let $\mathcal{B}$ be the $(d+1)$ -chain consisting of all the other $(d+1)$ -cells, then $\partial(\mathcal{A}) = \partial(\mathcal{B}) = \zeta_0$. Let v_1, v_2 be the two end vertices of $\theta(\sigma_\beta^{\mathcal{F}})$. Because the oppositely oriented simplex of $\vec{\sigma}_\beta^{\mathcal{F}}$ does not bound $\mathcal{V}$ in $\mathcal{Z}'$, it must be true that one of v_1, v_2 is in $\theta(\mathcal{A})$ and the other is in $\theta(\mathcal{B})$. We can let $(S, T) = (\theta(\mathcal{A}), \theta(\mathcal{B}))$ or $(\theta(\mathcal{B}), \theta(\mathcal{A}))$ based on which set contains the source of the flow network, then (S, T) is a cut of the flow network constructed in Algorithm 4.1. Furthermore, we have $\zeta_0 = \theta^{-1}(\xi(S, T))$ and $c(S, T) = w(\zeta_0) \leq w(\zeta)$. $\square$

The following theorem concludes the correctness of Algorithm 4.1:

Theorem 4.2. *Algorithm 4.1 computes a minimal persistent d -cycle for the given interval $[\beta, +\infty)$.*

Proof. First, the flow network (G, s_1, s_2) constructed by Algorithm 4.1 is valid. The reason is that, by Proposition 4.1, it cannot happen that the two oriented simplices of $\sigma_\beta^{\mathcal{F}}$ bound the same void of $\mathbb{R}^{d+1} \setminus |\tilde{K}|$. So $\sigma_\beta^{\mathcal{F}}$ must correspond to an edge of G . Then by Proposition 4.3 and 4.4, we can reach the conclusion. $\square$

5 Hardness for general complexes

Similar to the work [8], the NP-hardness proofs in this section accomplish the reduction with the help of a suspension operator. While Hatcher [21] defines this operator for general topological spaces, we need a definition of the operator for simplicial complexes and observe some of its properties that are useful for the proofs.

5.1 Suspension operator

Definition 5.1 (Suspension [20]). The *suspension* $\mathcal{S}K$ of a simplicial complex K is defined as a simplicial complex

$$\mathcal{S}K = \{\{\omega_1\}, \{\omega_2\}\} \cup K \cup \left(\bigcup_{\sigma \in K} \{\sigma \cup \{\omega_1\}, \sigma \cup \{\omega_2\}\} \right)$$

where ω_1, ω_2 are two extra vertices.

Remark 5.1. In the above definition, we denote a simplex by its set of vertices.

In the rest of this subsection, we let K be an arbitrary simplicial complex. Any simplex of the form $\sigma \cup \{\omega_i\}$ in $\mathcal{S}K$ is called a *suspended simplex*. The symbol $\mathcal{S}$ is also used to denote a linear map $\mathcal{S} : \mathbb{C}_q(K) \rightarrow \mathbb{C}_{q+1}(\mathcal{S}K)$, where $\mathcal{S}\sigma = \sigma \cup \{\omega_1\} + \sigma \cup \{\omega_2\}$ for any q -simplex σ of K . Note that since $\mathcal{S}$ is injective, the map $\mathcal{S}$ defines an isomorphism from $\mathbb{C}_q(K)$ to the image $\mathcal{S}(\mathbb{C}_q(K))$. For any chain $A \in \mathcal{S}(\mathbb{C}_q(K))$, we abuse the notation slightly by letting $\mathcal{S}^{-1}A$ denote the chain in $\mathbb{C}_q(K)$ mapped to A under $\mathcal{S}$.

Proposition 5.1. For any $q \geq 1$, the following diagram commutes:

$$\begin{array}{ccc} \mathbb{C}_q(K) & \xrightarrow{\partial} & \mathbb{C}_{q-1}(K) \\ \mathcal{S} \downarrow \approx & & \approx \downarrow \mathcal{S} \\ \mathcal{S}(\mathbb{C}_q(K)) & \xrightarrow{\partial} & \mathcal{S}(\mathbb{C}_{q-1}(K)) \end{array}$$

Proof. For any q -simplex $\sigma = \{v_0, \dots, v_q\}$ of K , we have

$$\begin{aligned} \partial(\mathcal{S}\sigma) &= \partial(\{v_0, \dots, v_q, \omega_1\} + \{v_0, \dots, v_q, \omega_2\}) \\ &= \sum_{i=0}^q \{v_0, \dots, \widehat{v}_i, \dots, v_q, \omega_1\} + \{v_0, \dots, v_q\} + \sum_{i=0}^q \{v_0, \dots, \widehat{v}_i, \dots, v_q, \omega_2\} + \{v_0, \dots, v_q\} \\ &= \sum_{i=0}^q (\{v_0, \dots, \widehat{v}_i, \dots, v_q, \omega_1\} + \{v_0, \dots, \widehat{v}_i, \dots, v_q, \omega_2\}) \\ &= \sum_{i=0}^q \mathcal{S}(\{v_0, \dots, \widehat{v}_i, \dots, v_q\}) = \mathcal{S}\left(\sum_{i=0}^q \{v_0, \dots, \widehat{v}_i, \dots, v_q\}\right) = \mathcal{S}\partial(\sigma) \end{aligned}$$

In the above equations, the notation $\widehat{v}_i$ means that v_i is deleted from the simplex. $\square$

Proposition 5.2. For $q \geq 1$ and any q -cycle ζ of $\mathcal{S}K$ containing only suspended simplices, one has $\zeta \in \mathcal{S}(\mathbb{C}_{q-1}(K))$.

Proof. For any suspended q -simplex $\sigma \cup \{\omega_i\}$ of ζ , if $\omega_i = \omega_1$, then $\sigma \cup \{\omega_2\}$ must also belong to ζ because no other suspended q -simplices of $\mathcal{S}K$ have σ in the boundary. If $\omega_i = \omega_2$, the same argument follows. $\square$

Proposition 5.3. If q is the top dimension of K and $q \geq 1$, then for any $A \in \mathbb{C}_{q+1}(\mathcal{S}K)$ such that $\partial(A)$ contains only suspended simplices, one has $A \in \mathcal{S}(\mathbb{C}_q(K))$.

Proof. Because q is the top dimension of K , A contains only suspended simplices. For any $\sigma \cup \{\omega_i\} \in A$, we have $\sigma \in \partial(\sigma \cup \{\omega_i\})$. If $\omega_i = \omega_1$, to make σ cancelled in $\partial(A)$, $\sigma \cup \{\omega_2\}$ must also belong to A because no other $(q+1)$ -simplices in $\mathcal{S}K$ have σ in the boundary. If $\omega_i = \omega_2$, the same argument follows. $\square$

5.2 Hardness for finite intervals

The following proposition helps to prove our conclusion of the hardness:

Proposition 5.4. *PCYC-FIN_{d-1} reduces to PCYC-FIN_d for $d \geq 2$.*

Proof. Given an instance $(K, \mathcal{F}, [\beta, \delta])$ of PCYC-FIN_{d-1}, where the i_{th} complex of $\mathcal{F}$ is denoted as K_i , we can assume the top dimension of K to be d . The reason is that if it were not, we can restrict $\mathcal{F}$ to the d -skeleton of K without affecting $D_{d-1}(\mathcal{F})$ and the persistent $(d-1)$ -cycles. Then, we let SK be the simplicial complex for the instance of PCYC-FIN_d we are going to construct. For any suspended d -simplex $\sigma \cup \{\omega_i\}$ of SK , let the weight of $\sigma \cup \{\omega_i\}$ be half of the weight of σ in K . Furthermore, let the weight of any non-suspended d -simplex of SK be the sum of all the weights of $(d-1)$ -simplices in K plus 1. We endow SK with a filtration $\mathcal{SF} : \emptyset = \hat{K}_0 \subseteq \hat{K}_1 \subseteq \dots \subseteq \hat{K}_{3n+2} = SK$, where n is the number of simplices of K . Denoting the i_{th} simplex added in $\mathcal{F}$ as σ_i and the i_{th} simplex added in $\mathcal{SF}$ as $\hat{\sigma}_i$, we let $\hat{\sigma}_1 = \{\omega_1\}$, $\hat{\sigma}_2 = \{\omega_2\}$, and for any $1 \leq i \leq n$, $\hat{\sigma}_{3i} = \sigma_i$, $\hat{\sigma}_{3i+1} = \sigma_i \cup \{\omega_1\}$, $\hat{\sigma}_{3i+2} = \sigma_i \cup \{\omega_2\}$.

We observe the following facts:

- (i) For any i , $\hat{\sigma}_{3i}$ is positive and pairs with $\hat{\sigma}_{3i+1}$ in $\mathcal{SF}$.
- (ii) For any i and j , if there is a $(d-1)$ -cycle created by σ_i which is a boundary in K_j , then there is a d -cycle created by $\hat{\sigma}_{3i+2}$ which is a boundary in $\hat{K}_{3j+2}$.
- (iii) For any i and j , if there is a d -cycle created by $\hat{\sigma}_{3i+2}$ which is a boundary in $\hat{K}_{3j+2}$, then there is a $(d-1)$ -cycle created by σ_i which is a boundary in K_j .

The correctness of (i) is not hard to verify. To verify (ii), we can suspend the $(d-1)$ -cycle and use Proposition 5.1 to reach the claim. The argument for (iii) is as follows: Consider a d -cycle $\hat{\zeta}_0$ created by $\hat{\sigma}_{3i+2}$ which is a boundary in $\hat{K}_{3j+2}$. For any non-suspended d -simplex σ of $\hat{\zeta}_0$, we add $\partial(\sigma \cup \{\omega_1\})$ to the cycle $\hat{\zeta}_0$ so that σ is canceled and only suspended simplices are added. Note that the adding process only adds d -simplices in $\hat{K}_{3i+2}$ and never cancels $\hat{\sigma}_{3i+2}$. After all non-suspended simplices of $\hat{\zeta}_0$ are canceled, we derive a d -cycle $\hat{\zeta}$ which is created by $\hat{\sigma}_{3i+2}$ and contains only suspended simplices. By Proposition 5.2, $\mathcal{S}^{-1}\hat{\zeta}$ is well defined. Since $\hat{\zeta}$ is homologous to $\hat{\zeta}_0$ in $\hat{K}_{3j+2}$, $\hat{\zeta}$ is also a boundary in $\hat{K}_{3j+2}$. Let $\hat{\zeta}$ be the boundary of a $(d+1)$ -chain $\hat{A}$ in $\hat{K}_{3j+2}$. Because $SK_j = \hat{K}_{3j+2}$, by Proposition 5.3, $\hat{A} \in \mathcal{S}(\mathcal{C}_d(K_j))$. Furthermore, by Proposition 5.1, we have $\mathcal{S}^{-1}\hat{\zeta} = \mathcal{S}^{-1}\partial(\hat{A}) = \partial(\mathcal{S}^{-1}\hat{A})$. So $\mathcal{S}^{-1}\hat{\zeta}$ is a $(d-1)$ -cycle created by σ_i which is a boundary in K_j .

From the above facts, it is immediate that $\hat{\sigma}_{3\beta+2}$ is a positive simplex in $\mathcal{SF}$ and pairs with $\hat{\sigma}_{3\delta+2}$ so that $[3\beta+2, 3\delta+2]$ is an interval in $D_d(\mathcal{SF})$. It is also true that there is a bijection from the persistent $(d-1)$ -cycles of $[\beta, \delta]$ to the persistent d -cycles of $[3\beta+2, 3\delta+2]$ containing only suspended simplices. Furthermore, the bijection preserves the weights of the cycles. From the weight assigning policy, the minimal persistent d -cycle of $[3\beta+2, 3\delta+2]$ must contain only suspended simplices, so this minimal persistent d -cycle of $[3\beta+2, 3\delta+2]$ induces a minimal persistent $(d-1)$ -cycle of $[\beta, \delta]$. Now we have reduced PCYC-FIN_{d-1} to PCYC-FIN_d. Furthermore, the reduction is in polynomial time and the size of $(SK, \mathcal{SF}, [3\beta+2, 3\delta+2])$ is a polynomial function of the size of $(K, \mathcal{F}, [\beta, \delta])$. $\square$

We have the following result from [13]:

Proposition 5.5. *PCYC-FIN₁ is NP-hard.*

Combining Proposition 5.4 and 5.5, we obtain the following theorem:

Theorem 5.1. *PCYC-FIN_d is NP-hard for $d \geq 1$.*

5.3 Hardness for infinite intervals

In this subsection, we prove that it is NP-hard to approximate WPCYC-INF_d with any fixed ratio. Let PROB be a minimization problem with solutions having positive costs. Given an instance $\mathcal{I}$ of PROB , let C^* be the cost of the minimal solution of $\mathcal{I}$. For $r \geq 1$, a solution of $\mathcal{I}$ with cost C is said to have an *approximation ratio* r if $C/C^* \leq r$ [10]. We let $\text{PROB}[r]$ denote the problem that asks for an approximate solution with ratio r given an instance of PROB . Moreover, in order to make approximation ratios well-defined for WPCYC-INF_d , we let WPCYC-INF_d^+ denote a subproblem of WPCYC-INF_d where all d -simplices are positively weighted.

Before proving the hardness result, we first recall the definition of the nearest codeword problem, which is NP-hard to approximate with any fixed ratio [8]:

Problem 5.1 (NR-CODE). Given an $l \times k$ full-rank matrix $\mathcal{A}$ over $\mathbb{Z}_2$ for $k < l$ and a vector $y_0 \in (\mathbb{Z}_2)^l \setminus \text{Img}(\mathcal{A})$, find a vector in $y_0 + \text{Img}(\mathcal{A})$ with the minimal Hamming weight.

Remark 5.2. The Hamming weight of a vector y , denoted as $\|y\|_H$, is the number of non-zero components in y .

Theorem 5.2. WPCYC-INF_2^+ is NP-hard to approximate with any fixed ratio.

Similar to the NP-hardness proof of homology localization in [8], our proof of Theorem 5.2 conducts the reduction from the NR-CODE problem. One may think that a direct reduction from homology localization may be more straightforward. However, such a reduction is not immediately evident. The two problems appear to be of different nature: While the homology localization problem asks for a minimal cycle in a given homology class, WPCYC-INF_2^+ asks for a minimal cycle in a complex containing a given simplex without referring to any particular homology class.

Proof. For any $r > 1$, we reduce the NP-hard problem $\text{NR-CODE}[2r]$ to $\text{WPCYC-INF}_2^+[r]$. Given an instance $(\mathcal{A}, y_0)$ of $\text{NR-CODE}[2r]$, we first compute the $(l - k) \times l$ parity check matrix $\mathcal{A}^\perp$ [8], which is a matrix such that $\text{Ker}(\mathcal{A}^\perp) = \text{Img}(\mathcal{A})$. Similar to the proof of Lemma 4.3.1 in [8], we then build a “tube complex” T_1 with $(l - k)$ 1-cells each of which is a 1-sphere and l 2-cells each of which is a 2-sphere with holes. The 2-cells of T_1 are attached to the 1-cells along the holes such that the boundary matrix ∂_2 of this tube complex equals $\mathcal{A}^\perp$. The “ q -chains” and “ q -cycles” for a tube complex are analogously defined as for a simplicial complex. We also assign a weight of 1 to each 2-cell of T_1 . By this construction, there is a straightforward bijection $\phi : (\mathbb{Z}_2)^l \rightarrow \text{C}_2(T_1)$, such that the Hamming weight of a vector equals the weight of the corresponding 2-chain. Note that $\text{Z}_2(T_1) = \text{Ker}(\partial_2) = \phi(\text{Ker}(\mathcal{A}^\perp)) = \phi(\text{Img}(\mathcal{A}))$. Let $\tilde{y}_0 = \phi(y_0)$, we then add a 2-cell $\hat{t}$ whose boundary equals $\partial_2(\tilde{y}_0)$ to T_1 and get a new tube complex T_2 . We call the 2-cycles in T_2 which are not in T_1 as the new 2-cycles in T_2 . Then $\hat{t} + \tilde{y}_0$ is a new 2-cycle in T_2 and the set of new 2-cycles in T_2 is $\hat{t} + \tilde{y}_0 + \text{Z}_2(T_1)$. We let the weight of $\hat{t}$ also be 1. Note that there is a bijection $\psi : y_0 + \text{Img}(\mathcal{A}) \rightarrow \hat{t} + \tilde{y}_0 + \text{Z}_2(T_1)$, where $\psi(y_0 + z) = \hat{t} + \tilde{y}_0 + \phi(z)$ for any $z \in \text{Img}(\mathcal{A})$, such that $w(\psi(y_0 + z)) = \|y_0 + z\|_H + w(\hat{t})$.

We then construct an instance of $\text{WPCYC-INF}_2^+[r]$ by first triangulating T_2 to get a simplicial complex K . We make K 2-weighted such that the sum of the weights of all triangles in any 2-cell of T_2 equals the weight of the 2-cell. It is not hard to make the size of K a polynomial function of the number of cells of T_2 . Let σ be a 2-simplex in the triangulation of the 2-cell $\hat{t}$. We build a filtration $\mathcal{F}$ of K with σ being the last simplex added. Let the index of σ in $\mathcal{F}$ be β . Then, $[\beta, +\infty)$ is an infinite interval of $\text{D}_2(\mathcal{F})$. Note that there is a bijection between the new 2-cycles in T_2 and the persistent 2-cycles of $[\beta, +\infty)$, where the weights of the cycles are preserved. Therefore, from the solution of $\text{WPCYC-INF}_2^+[r]$ with the input $(K, \mathcal{F}, [\beta, +\infty))$, we can derive a new 2-cycle

$\hat{t} + \tilde{y}_0 + \zeta$ of T_2 , where $\zeta \in Z_2(T_1)$ and $\hat{t} + \tilde{y}_0 + \zeta$ is an r -approximation of the minimal new 2-cycle. Let $\hat{t} + \tilde{y}_0 + \zeta^*$ be a minimal new 2-cycle of T_2 , we have

$$\frac{w(\hat{t} + \tilde{y}_0 + \zeta)}{w(\hat{t} + \tilde{y}_0 + \zeta^*)} \leq r \implies \frac{w(\hat{t}) + w(\tilde{y}_0 + \zeta)}{w(\hat{t}) + w(\tilde{y}_0 + \zeta^*)} \leq r \implies w(\tilde{y}_0 + \zeta) \leq r - 1 + rw(\tilde{y}_0 + \zeta^*)$$

We also have

$$1 \leq \frac{r}{r-1} w(\tilde{y}_0 + \zeta^*) \implies r - 1 \leq rw(\tilde{y}_0 + \zeta^*)$$

Therefore

$$w(\tilde{y}_0 + \zeta) \leq 2rw(\tilde{y}_0 + \zeta^*) \implies \|y_0 + \phi^{-1}(\zeta)\|_H \leq 2r\|y_0 + \phi^{-1}(\zeta^*)\|_H$$

Since $y_0 + \phi^{-1}(\zeta^*)$ is a minimal solution of $(\mathcal{A}, y_0)$, then $y_0 + \phi^{-1}(\zeta)$ is a $2r$ -approximation of the minimal solution of $(\mathcal{A}, y_0)$. Hence, we have reduced NR-CODE[$2r$] to WPCYC-INF $_2^+[r]$. Furthermore, the reduction is in polynomial time and the sizes of the instances are related by a polynomial function, so WPCYC-INF $_2^+[r]$ is NP-hard. $\square$

Theorem 5.3. *WPCYC-INF $_d^+$ is NP-hard to approximate with any fixed ratio for $d \geq 2$.*

Proof. For any $d \geq 3$ and $r \geq 1$, we reduce WPCYC-INF $_{d-1}^+[r]$ to WPCYC-INF $_d^+[r]$. Given an instance $(K, \mathcal{F}, [\beta, +\infty))$ of WPCYC-INF $_{d-1}^+[r]$, where the i th complex of $\mathcal{F}$ is denoted as K_i , let $K' = \mathcal{S}K_\beta^{d-1}$ where K_β^{d-1} is the $(d-1)$ -skeleton of K_β . We make K' d -weighted such that any d -simplex $\sigma \cup \{\omega_i\}$ of K' has half of the weight of σ in K . The complex K' is endowed with a filtration $\mathcal{F}'$ such that $\sigma_\beta^\mathcal{F} \cup \{\omega_2\}$ is the last simplex added to $\mathcal{F}'$. Let β' be the index of $\sigma_\beta^\mathcal{F} \cup \{\omega_2\}$ in $\mathcal{F}'$, then $[\beta', +\infty) \in D_d(\mathcal{F}')$. It is true that $\mathcal{S}$ restricts to a bijection from $Z_{d-1}(K_\beta)$ to $Z_d(K')$ preserving the weights of the cycles. Furthermore, for any $\zeta \in Z_{d-1}(K_\beta)$, ζ is a persistent $(d-1)$ -cycle of $[\beta, +\infty) \in D_{d-1}(\mathcal{F})$ if and only if $\mathcal{S}\zeta$ is a persistent d -cycle of $[\beta', +\infty) \in D_d(\mathcal{F}')$. Suppose that ζ' is a solution for the instance $(K', \mathcal{F}', [\beta', +\infty))$ of WPCYC-INF $_d^+[r]$, i.e., ζ' is an r -approximation of the minimal solution. Then, $\mathcal{S}^{-1}\zeta'$ is an r -approximation for the instance $(K, \mathcal{F}, [\beta, +\infty))$ of WPCYC-INF $_{d-1}^+[r]$. Therefore, the reduction is done. $\square$

6 Experimental results

We experiment with our algorithms for WPCYC-FIN $_2$ and WEPCYC-INF $_2$ on several volume datasets. Since volume data have a natural cubical complex structure, we adapt our implementation slightly in order to work on cubical complexes. The cubical complex for volume data consists of cells in dimensions from 0 to 3 with the underlying space homeomorphic to a 3-dimensional ball. Note that a filtration built from a volume dataset does not produce any infinite intervals. Hence, in order to test our algorithm for WEPCYC-INF $_2$, we take a finite interval and compute the minimal 2-cycle born at the birth time, which is exactly what WEPCYC-INF $_2$ computes. We use the **Gudhi** [29] library to build the filtrations and compute the persistence intervals. From the experiments, we can see that the minimal persistent 2-cycles computed by our algorithms capture various features of the data which originate from different fields. Note that the combustion, hurricane, and medical datasets are time-varying and we chose a single time frame to compute the persistent intervals and cycles.

Cosmology. The simulation data shown in Figure 4a from computational cosmology [2] consist of dark matter represented as particles. The thread-like structures in deep purple shown in Figure 4a correspond to sites of large scale structure formation. Galaxy clusters/superclusters are contained




Figure 4: (a,b) Cosmology dataset and the minimal persistent 2-cycles of the top five longest intervals. (c,d) Turbulent combustion dataset and its corresponding minimal persistent 2-cycles.




Figure 5: (a,b) Minimal persistent 2-cycles for the hurricane model. (c) Minimal persistent 2-cycles of the larger intervals for the human skull. i: Right and left cheek muscles with the right one rotated for better visibility. ii: Right and left eyes. iii: Jawbone. iv: Nose cartilage. v: Nerves in the parietal lobe.

in such large scale structures. Figure 4b shows the minimal persistent 2-cycles of the top five longest intervals computed by our algorithms and these cycles precisely represent the top five galaxy clusters/superclusters in volume.

Combustion. The data shown in Figure 4c correspond to the physical variable^{||} χ from a model of a turbulent combustion process. The variable χ represents scalar dissipation rate and provides a measure of the maximum possible chemical reaction rate. The minimal persistent 2-cycles shown in Figure 4d represent areas with high value of χ .

Hurricane. This dataset^{**} with 11 physical variables corresponds to the devastating hurricane named Isabel. We down-sampled the data into a resolution of $250 \times 250 \times 50$ and worked with

^{||} A physical variable defines a scalar value of a certain kind on each point.

^{**} The Hurricane Isabel data is produced by the Weather Research and Forecast (WRF) model, courtesy of NCAR, and the U.S. National Science Foundation (NSF).




Figure 6: (a) Cubic lattice structure of $BaTiO_3$ (courtesy Springer Materials [27]) with diffused structure in backdrop. (b) Minimal persistent 2-cycles computed on the original function. (c) Minimal persistent 2-cycles computed on the negated function. (d) Minimal persistent 2-cycles computed on the negated function of a tetragonal lattice structure of $BaTiO_3$. The inlaid picture [27] illustrates the bonds of the structure.

two physical variables. The minimal persistent 2-cycle colored blue in Figure 5a is computed on the cloud-volume variable and extracts the eye of the hurricane. The minimal persistent 2-cycle colored green in Figure 5b is computed on the pressure variable and captures the jagged shape of the pressure variation around the hurricane.

Medical imaging. This dataset from the ADNI [26] project contains the MRI scan of a healthy human skull. The minimal persistent 2-cycles corresponding to the larger intervals as shown in Figure 5c are computed from two time frames. They extract significant features such as eyes, cartilages, nerves, and muscles.

Material science. We consider the atomic configuration of $BaTiO_3$, which is a ferroelectric material used for making capacitors, transducers, and microphones. Figure 6a shows the atomic configuration of the molecule, where the red, grey, and green balls denote the Oxygen, Titanium, and Barium atoms separately and the radii of the balls equal the radii of the corresponding atoms. Volume data are built by uniformly sampling a $3 \times 3 \times 3$ lattice structure similar to the one shown in Figure 6a, with the step width equal to one *angstrom* (note that Figure 6a only shows a $2 \times 2 \times 2$ lattice structure). Scalar value on a point of the volume is determined as follows: For each atom, let the distance from the point to the atom's center be d , then the scalar value of the point contributed by the atom is $\max\{w(r - d)/r, 0\}$, where r is the radius of the atom and w is the atomic weight. The scalar value on the point is then equal to the sum of the above values contributed by all atoms. For the purpose of this experiment, we computed minimal persistent 2-cycles on both the original scalar function and its negated one. Figure 6b shows a portion of the minimal persistent 2-cycles computed on the original function, where the purple, red, and green cycles correspond to atoms of Barium, Titanium, and Oxygen respectively. In our experiment, every atom corresponds to such a minimal persistent 2-cycle of a long interval. Figure 6c shows a portion of the minimal persistent 2-cycles computed on the negated function, where the cycles complement the Barium atoms. Figure 6d shows the output on the negated function from a tetragonal lattice structure [27], where the atomic bonds are not straight (see Figure 6d inlay). The stretch on the lattice structure leads to minimal persistent 2-cycles with non-trivial genus.

7 Conclusions

In this paper, we inspect the computational complexity for several problems concerning minimal persistent cycles. We expand the hardness results found in [13] and discover the cases that are NP-hard and others that are solvable in polynomial time. For general complexes, we conclude that the computation is NP-hard over all dimensions for finite intervals and NP-hard over dimension greater than one for infinite intervals. Besides, we find the problems to be tractable in dimension d if the given complex is a weak $(d + 1)$ -pseudomanifold and, for infinite intervals, if the weak $(d + 1)$ -pseudomanifold is embedded in $\mathbb{R}^{d+1}$.

This research leads to some open questions concerning persistent cycles:

i. In our experiments, some persistent cycles correspond to important features of the data (see Section 6). However, we also ran into some intervals whose persistent cycles do not have obvious meanings. If there are ways to design filtrations for data such that persistent cycles are related to the important features, then the prospect for the application of persistent cycles or persistence in general would be more extensive.

ii. As found in [13], persistent cycles are not stable in general even when only the weights of the cycles are considered. It will be helpful to figure out assumptions that are still relevant in practice, but under which the persistent cycles remain stable.

iii. We have presented $O(n^2)$ -time algorithms for computing a minimal persistent cycle for a given interval. A natural question is whether this time complexity can be improved. Furthermore, can we devise a better algorithm to compute minimal persistent cycles for all intervals (i.e., the *minimal persistent basis* [13]), improving upon the obvious $O(n^3)$ -time algorithm that runs our algorithms on each interval?

Acknowledgments:

This research was conducted with the support of the NSF grants CCF-1740761 and CCF-1839252. We thank the anonymous reviewers for insightful comments.

A Proof of Theorem 4.1

We first define some symbols used in this section. The interior of a set U is denoted by $\text{Int}(U)$. The boundary of a topological ball $\mathbb{B}$ is denoted by $\text{bd}(\mathbb{B})$. The set of q -cofaces of a simplex σ in a Δ -complex [21] K is denoted by $\text{cof}_q^K(\sigma)$.

The proof of Theorem 4.1 is based on the extended Jordan–Brouwer separation theorem (Theorem A.1) by Alexander [1]. The statement of the theorem depends on the following definition:

Definition A.1 (Pseudomanifold). A simplicial complex K is a q -pseudomanifold if K is a pure q -complex and each $(q - 1)$ -simplex is a face of exactly two q -simplices in K .

Remark A.1. Note that definitions for q -pseudomanifolds, such as in [28], typically assume the complex to be q -connected.

Theorem A.1. Let $q > 1$ and $\mathcal{M}$ be a finite $(q - 1)$ -connected $(q - 1)$ -pseudomanifold embedded in $\mathbb{R}^q$, then $\mathbb{R}^q \setminus |\mathcal{M}|$ has exactly 2 connected components.

Now we can finish our proof:




Figure 7: (a) An example of the “de-contraction” of σ^d for $d = 1$, where a 1-simplex in the left simplicial complex turns into two curved 1-simplices with identical boundary in the right Δ -complex. The topological 2-ball $\mathbb{B}'$ is the one bounded by the two curved 1-simplices. (b) Left to middle: An example demonstrating the void boundary correspondence from $\tilde{K}$ to $\tilde{K}'$ for $d = 1$. After a 1-simplex is de-contracted, the shaded void for $\tilde{K}$ corresponds to the shaded void for $\tilde{K}'$ and their boundaries (dashed line) can be identified. Middle to right: The “de-pinching” properly separates apart incident edges (1-simplices) for the two vertices (0-simplices) having more than two 1-cofaces. The complex $\mathcal{M}_h$ (on the right) then becomes a pseudomanifold. To deform $\mathcal{M}_h$ back to $\mathcal{M}$ (in this example $\mathcal{M} = \tilde{K}'$), only points in $\mathcal{B}$ (unshaded region) are contracted.

Proof of Theorem 4.1. The general idea of the proof is as follows: Using a trick which we call the “de-contracting”, we first create a Δ -complex $\tilde{K}'$ where each oriented simplex of $\vec{\zeta}_j$ uniquely corresponds to an unoriented simplex. Then, using a trick which we call the “de-pinching”, we show that $\vec{\zeta}_j$ is the boundary of a region $\mathcal{A}$. Finally, from the above fact, we use proof by contradiction to reach the conclusion. Figure 7b gives an example of the “de-contracting” and “de-pinching”.

First, let Σ' be the set of d -simplices of $\tilde{K}$ whose both oriented simplices are in $\vec{\zeta}_j$. For a d -simplex σ^d of Σ' , we can let $\mathbb{B}'$ be a topological $(d+1)$ -ball residing in $\mathbb{R}^{d+1}$ such that $\text{bd}(\mathbb{B}')$ equals two d -simplices with boundaries glued together. We then homeomorphically map points of $\mathbb{R}^{d+1} \setminus \sigma^d$ to $\mathbb{R}^{d+1} \setminus \mathbb{B}'$. By taking care of the mapping near the boundary of $\mathbb{B}'$, we can get a new ambient $\mathbb{R}^{d+1}$ and a new Δ -complex where all simplices of $\tilde{K}$ are untouched except that σ^d now corresponds to the two d -simplices bounding $\mathbb{B}'$. We can also think of the above process as “de-contracting” the topological d -ball σ^d into the topological $(d+1)$ -ball $\mathbb{B}'$ so that σ^d turns into two separate d -simplices with identical $(d-1)$ -faces (see Figure 7a for an example). After doing the “de-contraction” for all d -simplices in Σ' , we get a Δ -complex $\tilde{K}'$. It is true that an oriented boundary d -simplex in $\tilde{K}$ can be naturally identified as an oriented boundary d -simplex in $\tilde{K}'$. It is also true that the groups of oriented boundary d -simplices in $\tilde{K}$ are still groups of oriented boundary d -simplices in $\tilde{K}'$ under the natural identification. So we can let $\vec{\zeta}_j$ denote the same group of oriented d -simplices in $\tilde{K}'$. The construction guarantees that if $\vec{\zeta}_j$ is the boundary of a void of $\mathbb{R}^{d+1} \setminus |\tilde{K}'|$, then $\vec{\zeta}_j$ is also the boundary of a void of $\mathbb{R}^{d+1} \setminus |\tilde{K}|$. So we only need to show that $\vec{\zeta}_j$ is the boundary of a void of $\mathbb{R}^{d+1} \setminus |\tilde{K}'|$ (see Figure 7b for an example). From now on, we always treat $\vec{\zeta}_j$ as a set of oriented d -simplices as well as a d -cycle (with $\mathbb{Z}$ coefficients) in $\tilde{K}'$.

Since different oriented simplices of $\vec{\zeta}_j$ correspond to different unoriented simplices in $\tilde{K}'$, we define a bijection $\psi : \vec{\zeta}_j \rightarrow \zeta$. The bijection ψ maps each oriented simplex of $\vec{\zeta}_j$ to its corresponding unoriented simplex and ζ is the image of this mapping. We then let $\mathcal{M}$ be the closure of the simplicial set ζ . Note that ζ is a d -cycle (with $\mathbb{Z}_2$ coefficients) of $\tilde{K}'$ and $\mathcal{M}$ is a subcomplex of $\tilde{K}'$. Therefore, each $(d-1)$ -simplex is a face of an even number of d -simplices in $\mathcal{M}$. We first pick a $(d-1)$ -simplex σ^{d-1} of $\mathcal{M}$ such that $|\text{cof}_d^{\mathcal{M}}(\sigma^{d-1})| > 2$, then pick two d -simplices σ_0^d and σ_1^d from $\text{cof}_d^{\mathcal{M}}(\sigma^{d-1})$ such that $\psi^{-1}(\sigma_0^d)$ and $\psi^{-1}(\sigma_1^d)$ are paired in the void boundary reconstruction for $\tilde{K}'$.

It is then true that $\sigma_0^d \cup \sigma_1^d$ forms a topological d -ball $\mathbb{B}_1^d$ containing σ^{d-1} . Forming the topological d -balls for all such pairs of d -simplices in $\text{cof}_d^{\mathcal{M}}(\sigma^{d-1})$, we get a set of d -balls $\{\mathbb{B}_1^d, \dots, \mathbb{B}_\kappa^d\}$ for $\kappa = |\text{cof}_d^{\mathcal{M}}(\sigma^{d-1})|/2$. For each i , we slightly move $\mathbb{B}_i^d \setminus \text{Int}(\sigma^{d-1})$ while keeping $\text{bd}(\mathbb{B}_i^d)$ untouched. We then take the closure of each $\mathbb{B}_i^d \setminus \text{Int}(\sigma^{d-1})$ to get a new Δ -complex $\mathcal{M}_1$ in which the $\mathbb{B}_i^d$'s have their interiors disjoint. Note that in $\mathcal{M}_1$, σ^{d-1} now corresponds to κ different $(d-1)$ -simplices sharing the boundary. We can repeat the above “de-pinching” process for each $(d-1)$ -simplex having more than two d -cofaces in $\mathcal{M}$ and then get a sequence of Δ -complexes $(\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_h)$. In the sequence, $\mathcal{M}_0 = \mathcal{M}$ and $\mathcal{M}_i$ is derived from $\mathcal{M}_{i-1}$ by doing the “de-pinching” on a $(d-1)$ -simplex. It is then true that $\mathcal{M}_h$ is a pure d -dimensional d -connected Δ -complex where each $(d-1)$ -simplex is a face of exactly two d -simplices. Since we can subdivide $\mathcal{M}_h$ to make it a simplicial complex, by Theorem A.1, $|\mathcal{M}_h|$ must separate $\mathbb{R}^{d+1}$ into two connected components. Note that for each i , we can treat $\mathbb{R}^{d+1} \setminus |\mathcal{M}_i|$ as a subset of $\mathbb{R}^{d+1} \setminus |\mathcal{M}_{i+1}|$ because to deform $\mathcal{M}_{i+1}$ back to $\mathcal{M}_i$, we only need to contract some points in $\mathbb{R}^{d+1} \setminus |\mathcal{M}_{i+1}|$ to points in $|\mathcal{M}_{i+1}|$. Then the connected components of $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$ are still connected in $\mathbb{R}^{d+1} \setminus |\mathcal{M}_h|$. Since all oriented d -simplices of $\vec{\zeta}_j$ bound the same void of $\mathbb{R}^{d+1} \setminus |\tilde{K}'|$, we can let this void be $\mathcal{V}$. The void $\mathcal{V}$ is still connected in $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$ because $\mathbb{R}^{d+1} \setminus |\tilde{K}'| \subseteq \mathbb{R}^{d+1} \setminus |\mathcal{M}|$. Therefore, $\mathcal{V}$ is still connected in $\mathbb{R}^{d+1} \setminus |\mathcal{M}_h|$. We can let $\mathcal{A}$ be the connected component of $\mathbb{R}^{d+1} \setminus |\mathcal{M}_h|$ containing $\mathcal{V}$ and let $\mathcal{B}$ be the other connected component. The d -simplices in $\mathcal{M}$ and $\mathcal{M}_h$ can be identified because going from each $\mathcal{M}_i$ to $\mathcal{M}_{i+1}$ the interior of each d -simplex is never touched. Therefore, ζ is still a d -cycle (with $\mathbb{Z}_2$ coefficients) in $\mathcal{M}_h$. We then have that the two d -cycles (with $\mathbb{Z}$ coefficients) in $\mathcal{M}_h$, which are derived from the two consistent orientations of simplices of ζ , bound $\mathcal{A}$ and $\mathcal{B}$. Then, as one of the two d -cycles (with $\mathbb{Z}$ coefficients) derived from ζ , $\vec{\zeta}_j$ must be the boundary of $\mathcal{A}$ or $\mathcal{B}$ in $\mathcal{M}_h$. We have that $\vec{\zeta}_j$ bounds $\mathcal{A}$ because $\mathcal{B}$ does not contain points from $\mathcal{V}$. A fact about our construction is that to deform each $\mathcal{M}_i$ back into $\mathcal{M}_{i-1}$, we only need to contract points in $\mathcal{B}$. This implies that $\mathcal{A}$ is still a void of $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$ with boundary $\vec{\zeta}_j$ (see Figure 7b for an example).

To prove that $\vec{\zeta}_j$ is the boundary of a void of $\mathbb{R}^{d+1} \setminus |\tilde{K}'|$, we only need to show that there are no oriented d -simplices which are in the boundary of $\mathcal{V}$ but do not belong to $\vec{\zeta}_j$. For contradiction, suppose that there is such an oriented d -simplex $\vec{\sigma}^d$. Then $\vec{\sigma}^d$ must not be oppositely oriented to any oriented simplex of $\vec{\zeta}_j$ because otherwise $\vec{\sigma}^d$ would bound another connected component of $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$ and thus bound another connected component of $\mathbb{R}^{d+1} \setminus |\tilde{K}'|$. Let σ^d be the unoriented d -simplex of $\vec{\sigma}^d$, then $\sigma^d \notin \mathcal{M}$ because otherwise $\vec{\sigma}^d$ would be oppositely oriented to an oriented simplex of $\vec{\zeta}_j$. Since $\sigma^d \notin \mathcal{M}$, the interior of σ^d must reside in $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$. From now on, we always treat $\mathcal{A}$ as a void of $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$. Then among all voids of $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$, the interior of σ^d resides in $\mathcal{A}$. This is because $\mathcal{A}$ is the void of $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$ containing $\mathcal{V}$. If σ^d resides in a void other than $\mathcal{A}$, points to either side of σ^d cannot be from $\mathcal{V}$. Since $\tilde{K}'$ is d -connected, there must be a sequence of d -simplices $(\sigma_0^d, \dots, \sigma_l^d)$ of $\tilde{K}'$ such that $\sigma_0^d = \sigma^d$, $\sigma_l^d \in \mathcal{M}$, and $\sigma_i^d, \sigma_{i+1}^d$ share a $(d-1)$ -face for each i such that $0 \leq i < l$. Because the interior of σ_l^d is not in $\mathcal{A}$, we can let $\sigma_{l'}^d$ be the first d -simplex in the sequence whose interior is not in $\mathcal{A}$, then $l' \neq 0$ and the interior of $\sigma_{l'-1}^d$ is in $\mathcal{A}$. Let $\sigma_{l'-1}^{d-1}$ be the $(d-1)$ -face shared by $\sigma_{l'-1}^d$ and $\sigma_{l'}^d$, we claim that $\sigma_{l'-1}^{d-1} \in \mathcal{M}$. If $\sigma_{l'}^d \in \mathcal{M}$, then it is obvious that $\sigma_{l'-1}^{d-1} \in \mathcal{M}$. If $\sigma_{l'}^d \notin \mathcal{M}$, then it is also true that $\sigma_{l'-1}^{d-1} \in \mathcal{M}$ because otherwise the interiors of $\sigma_{l'-1}^d$ and $\sigma_{l'}^d$ would be connected in $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$. Around the neighborhood of $\sigma_{l'-1}^{d-1}$ during the void boundary reconstruction for $\tilde{K}'$, any two paired oriented simplices from $\vec{\zeta}_j$ enclose a region residing in $\mathcal{A}$. Because of the nature of the pairing, $\sigma_{l'-1}^d$ cannot be contained in any of the regions enclosed by the paired oriented simplices from $\vec{\zeta}_j$. Since $\vec{\zeta}_j$ is the boundary of the void $\mathcal{A}$ of $\mathbb{R}^{d+1} \setminus |\mathcal{M}|$, all other regions in the neighborhood of $\sigma_{l'-1}^{d-1}$ must not be in $\mathcal{A}$. This implies that $\sigma_{l'-1}^d$ is not in $\mathcal{A}$, which is a contradiction. $\square$

References

- [1] James W. Alexander. A proof and extension of the Jordan-Brouwer separation theorem. *Transactions of the American Mathematical Society*, 23(4):333–349, 1922.
- [2] Ann S. Almgren, John B. Bell, Mike J. Lijewski, Zarija Lukić, and Ethan Van Andel. Nyx: A massively parallel AMR code for computational cosmology. *The Astrophysical Journal*, 765(1):39, feb 2013.
- [3] Glencora Borradaile, Erin Wolf Chambers, Kyle Fox, and Amir Nayyeriy. Minimum cycle and homology bases of surface-embedded graphs. *Journal of Computational Geometry*, 8(2), 2017.
- [4] Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Minimum cuts and shortest homologous cycles. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 377–385. ACM, 2009.
- [5] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 237–246. ACM, 2009.
- [6] Frédéric Chazal, Vin De Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. Springer, 2016.
- [7] Chao Chen and Daniel Freedman. Measuring and computing natural generators for homology groups. *Computational Geometry*, 43(2):169–181, 2010.
- [8] Chao Chen and Daniel Freedman. Hardness results for homology localization. *Discrete & Computational Geometry*, 45(3):425–448, 2011.
- [9] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
- [10] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. Introduction to algorithms, 3rd edition. chapter 35, Approximation Algorithms. MIT Press, 2009.
- [11] Tamal K. Dey, Fengtao Fan, and Yusu Wang. Computing topological persistence for simplicial maps. In *Proceedings of the thirtieth annual symposium on Computational geometry*, page 345. ACM, 2014.
- [12] Tamal K. Dey, Anil N. Hirani, and Bala Krishnamoorthy. Optimal homologous cycles, total unimodularity, and linear programming. *SIAM Journal on Computing*, 40(4):1026–1044, 2011.
- [13] Tamal K. Dey, Tao Hou, and Sayan Mandal. Persistent 1-cycles: Definition, computation, and its application. In *Computational Topology in Image Context - 7th International Workshop, CTIC 2019, Málaga, Spain, January 24-25, 2019, Proceedings*, pages 123–136, 2019.
- [14] Tamal K. Dey, Jian Sun, and Yusu Wang. Approximating loops in a shortest homology basis from point data. In *Proceedings of the twenty-sixth annual symposium on Computational geometry*, pages 166–175. ACM, 2010.
- [15] Herbert Edelsbrunner and John Harer. *Computational topology: an introduction*. American Mathematical Soc., 2010.

- [16] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 454–463. IEEE, 2000.
- [17] Kevin Emmett, Benjamin Schweinhart, and Raul Rabadan. Multiscale topology of chromatin folding. In *Proceedings of the 9th EAI international conference on bio-inspired information and communications technologies (formerly BIONETICS)*, pages 177–180. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2016.
- [18] Jeff Erickson and Kim Whittlesey. Greedy optimal homotopy and homology generators. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1038–1046. Society for Industrial and Applied Mathematics, 2005.
- [19] Emerson G. Escolar and Yasuaki Hiraoka. Optimal cycles for persistent homology via linear programming. In *Optimization in the Real World*, pages 79–96. Springer, 2016.
- [20] Davide L. Ferrario and Renzo A. Piccinini. *Simplicial structures in topology*. Springer Science & Business Media, 2010.
- [21] A. Hatcher. *Algebraic topology*. Cambridge University Press, 2002.
- [22] John Lee. *Introduction to topological manifolds*, volume 202. Springer Science & Business Media, 2010.
- [23] James R. Munkres. *Elements of algebraic topology*. CRC Press, 2018.
- [24] Ippei Obayashi. Volume-optimal cycle: Tightest representative cycle of a generator in persistent homology. *SIAM Journal on Applied Algebra and Geometry*, 2(4):508–534, 2018.
- [25] James B. Orlin. Max flows in $O(nm)$ time, or better. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 765–774. ACM, 2013.
- [26] Ronald Carl Petersen, PS Aisen, Laurel A Beckett, MC Donohue, AC Gamst, Danielle J Harvey, CR Jack, WJ Jagust, LM Shaw, AW Toga, et al. Alzheimer’s disease neuroimaging initiative (ADNI): clinical characterization. *Neurology*, 74(3):201–209, 2010.
- [27] Pierre Villars (Chief Editor). PAULING FILE in: Inorganic Solid Phases. SpringerMaterials (online database), Springer, Heidelberg (ed.).
- [28] Edwin H. Spanier. *Algebraic topology*, volume 55. Springer Science & Business Media, 1989.
- [29] The GUDHI Project. *GUDHI User and Reference Manual*. GUDHI Editorial Board, 2015.
- [30] Pengxiang Wu, Chao Chen, Yusu Wang, Shaoting Zhang, Changhe Yuan, Zhen Qian, Dimitris Metaxas, and Leon Axel. Optimal topological cycles and their application in cardiac trabeculae restoration. In *International Conference on Information Processing in Medical Imaging*, pages 80–92. Springer, 2017.
- [31] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.